Fast computation of all maximum acyclic agreement forests for two rooted binary phylogenetic trees

Benjamin Albrecht *

Institut für Informatik, Ludwig-Maximilians-Universität, Germany

arXiv:1512.05656v1 [q-bio.PE] 17 Dec 2015

Abstract. Evolutionary scenarios displaying reticulation events are often represented by rooted phylogenetic networks. Due to biological reasons, those events occur very rarely, and, thus, networks containing a minimum number of such events, so-called minimum hybridization networks, are of particular interest for research. Moreover, to study reticulate evolution, biologist need not only a subset but *all* of those networks. To achieve this goal, the less complex concept of rooted phylogenetic trees can be used as building block. Here, as a first important step, the trees are disjoint into common parts, so-called maximum acyclic agreement forests, which can then be turned into minimum hybridization networks by applying further network building algorithms. In this paper, we present two modifications of the first non-naive algorithm — called allMAAFs — computing *all* maximum acyclic agreement forests for two rooted binary phylogenetic trees on the same set of taxa. By a simulation study, we indicate that through these modifications the algorithm is on average 8 times faster than the original algorithm making this algorithm accessible to larger input trees and, thus, to a wider range of biological problems.

Keywords Directed acyclic graphs · Hybridization · Maximum acyclic agreement forests · Bounded search · Phylogenetics

* Benjamin Albrecht
Institut für Informatik, Ludwig-Maximilians-Universität, Germany
Tel.: +49-89-2180-4069
E-mail: albrecht@bio.ifi.lmu.de

1 Introduction

Phylogenetic studies deal with the reconstruction of evolutionary histories, which are often represented by rooted phylogenetic trees in which each node represents a certain speciation event. The representation of reticulation events combining different evolutionary histories, however, requires the more general concept of phylogenetic networks, which are more complex structures that can additionally consist of nodes of in-degree larger than one. Hybridization, for instance, is an example for such a reticulation event merging a sizable percentage of two genomes of two different species. Now, due to those reticulation events, for a set of species there can exist incongruent gene trees, which are phylogenetic trees based on different genes sampled for this particular set of species all representing different evolutionary histories. On should, however, keep in mind that there can also exist other biological effects leading to incongruent gene trees, e.g. incomplete lineage sorting, still allowing to model evolutionary histories by phylogenetic trees.

As it is the case for hybridization events, a reticulation event is in general expected to occur very rarely. Thus, in order to study reticulate evolution, one is interested in the minimum number of such events explaining a set of incongruent gene trees. A further more complicated but strongly connected problem is the computation of *rooted phylogenetic networks* displaying evolutionary scenarios containing such events. In such a network each node with in-degree one reflects a putative speciation event, whereas each node of in-degree of at least two, a so-called reticulation node (or, in respect to hybridization, hybridization node) represents a putative reticulation event.

Phylogenetic studies are often joint works between biologists, mathematicians and computer scientists; usually biologists produce certain data, mathematicians develop a model for analyzing this data and computer scientists implement programs, which enable the investigation of the data based on these models. Regarding the investigation of reticulate evolution, the underlying problem is in general highly combinatorial, which complicates the work of computer scientists as naive approaches are typically insufficient for analyzing biologically relevant data.

In this work, we tackle this problem by focusing on the first non-naive algorithm allMAAFs [12] that enables the computation of *all* maximum acyclic agreement forests for two rooted binary phylogenetic trees, which can be considered, from a mathematical point of view, as an intermediate step in calculating *all* possible rooted phylogenetic networks. This is the case, since those networks can then be computed by gluing the components of such maximum acyclic agreement forests back together in a specific way which, for example, has already been demonstrated by the algorithm HybridPhylogeny [4]. The computation of a maximum acyclic agreement forest for two binary phylogenetic trees, however, is a well-known *NP-hard* problem [5] and, thus, in order to apply this step to large input trees and thereby making the algorithm accessible to a wider range of biological problems, we have worked out two modifications significantly improving the practical running time of the algorithm allMAAFs.

Notice that so far there have been developed several approaches for analyzing reticulation events [1, 7, 8, 13]. Whereas some of these approaches just enable the computation of hybridization numbers, some additionally can be used to generate a set of minimum hybridization networks each representing a particular evolutionary scenario. Notice, however, that until now none of these methods is able to calculate *all* minimum hybridization networks, which, as already emphasized, under a biological point of view, is an important feature for studying reticulate evolution.

The paper is organized as follows. First, we give all basic notations and some well-known results of phylogenetics research, which are used throughout this work. Next, we describe two modifications allMAAFs^1 and allMAAFs^2 of the algorithm allMAAFs by presenting its respective pseudo code. In a subsequent step, the correctness of each of these modified algorithms is established by formal proofs and, finally, in order to indicate the speedup obtained from applying these modifications, we present the results of a simulation study comparing runtimes attained from its respective implementations. Lastly, we finish this paper by briefly discussing the theoretical worst-case running time of both algorithms allMAAFs^1 and allMAAFs^2.

2 Preliminaries

In this section, we give all formal definitions that are used throughout this work for describing and discussing the original algorithm allMAAFs as well as our two modified algorithms allMAAFs[1] and allMAAFs[2]. These definitions correspond to those given in the work of Huson *et. al.* [10] and Scornavacca *et. al.* [12]. We assume here that the reader is familiar with general graph-theoretic concepts.

Phylogenetic trees. A *rooted phylogenetic $\mathcal{X}$-tree* T is a directed acyclic connected graph whose edges are directed from the root to the leaves as defined in the following. There is exactly one node of in-degree 0, namely the *root* of T. The set of nodes of out-degree 0 is called the *leaf set of* T and is labeled one-to-one by a *taxa set* $\mathcal{X}$, also denoted by $\mathcal{L}(T)$. Here, the taxa set $\mathcal{X}$ usually consists of certain species or genes whose evolution is outlined by T. The tree T is called *binary* if all of its nodes, except the root, provide an in-degree of 1 and if all of its nodes, except all leaves (the so-called *inner or internal nodes*) provide an out-degree of 2. Given a node v, the label set $\mathcal{L}(v)$ refers to all taxa that are contained in the subtree rooted at v. In addition, given a set of rooted phylogenetic $\mathcal{X}$-trees $\mathcal{F}$, by $\mathcal{L}(\mathcal{F})$ we refer to the union of each leaf set of each tree in $\mathcal{F}$.

Now, given a set of taxa $\mathcal{X}' \subseteq \mathcal{L}$, the notation $T(\mathcal{X}')$ refers to the minimal connected subgraph of T containing $\mathcal{X}'$. A *restriction* of T to $\mathcal{X}'$, shortly denoted by $T|_{\mathcal{X}'}$, is a rooted phylogenetic tree that is obtained from $T(\mathcal{X}')$ by suppressing each node of both in- and out-degree 1. Moreover, given two phylogenetic trees T_1 and T_2 in which $\mathcal{X}_1$ and $\mathcal{X}_2$, with $\mathcal{X}_1 \subseteq \mathcal{X}_2$, denotes the taxa set of T_1 and T_2, respectively, we say T_1 is *contained* in T_2, shortly denoted by $T_1 \subseteq T_2$, if $T_2|_{\mathcal{X}_1}$ is isomorphic to T_1.

Next, the *lowest common ancestor* of a phylogenetic $\mathcal{X}$-tree T in terms of a taxa set $\mathcal{X}' \subseteq \mathcal{X}$ is the node v in T with $\mathcal{X}' \subseteq \mathcal{L}(v)$ such that there does not exist another node v' in T with $\mathcal{X}' \subseteq \mathcal{L}(v')$ and $\mathcal{L}(v') \subseteq \mathcal{L}(v)$. In the following, such a node is shortly denoted by $\mathrm{LCA}_T(\mathcal{X}')$.

Lastly, given a set of phylogenetic trees $\mathcal{F}$ as well as an edge set E' that is contained in $\mathcal{F}$ such that for each pair of edges $e_1, e_2 \in E'$, with $e_1 \neq e_2$, e_1 is not adjacent to e_2. Then, by $\mathcal{F} - E'$ we refer to the set $\mathcal{F}'$ of trees that is obtained from $\mathcal{F}$ by *cutting* E'. More precisely, first each edge in E' is deleted in $\mathcal{F}$ and then each node of both in- and out-degree 1 is suppressed. Notice that by deleting an edge of a tree F in $\mathcal{F}$, this tree is separated into two parts F_a and F_b so that the resulting forest equals $\mathcal{F} \setminus \{F\} \cup \{F_a, F_b\}$. Consequently, after deleting E' from $\mathcal{F}$ the resulting set $\mathcal{F}'$ contains precisely $|\mathcal{F}| + |E'|$ trees.

During the algorithm allMAAFs, an agreement forests for two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2 is calculated by cutting down one of both trees, say T_2, into several components such that each component corresponds to a restricted subtree of T_1. In order to keep track of the component F_ρ containing the root, the root of T_2 has to be a node that is marked in a specific way. Thus, throughout this paper, we regard the root of each rooted binary phylogenetic $\mathcal{X}$-tree as a node that is attached to its original root and to a taxon $\rho \notin \mathcal{X}$ (cf. Fig. 1(a)).

Phylogenetic networks. A *rooted phylogenetic network* N on $\mathcal{X}$ is a rooted connected digraph whose edges are directed from the root to the leaves as defined in the following. There is exactly one node of in-degree 0, namely the *root*, and no nodes of both in- and out-degree 1. The set of nodes of out-degree 0 is called the *leaf set of* N and is labeled one-to-one by the *taxa set* $\mathcal{X}$, also denoted by $\mathcal{L}(N)$. In contrast to a phylogenetic tree, such a network may contain undirected but not any directed cycles. Consequently, N can contain nodes of in-degree larger than or equal to 2, which are called *reticulation nodes* or *hybridization nodes*. Moreover, all edges that are directed into such a reticulation node are called *reticulation edges* or *hybridization edges*.

Hybridization Networks. A hybridization network N for two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2 is a rooted phylogenetic network on $\mathcal{X}$ *displaying* T_1 and T_2. More precisely, this means that for each tree $T \in \{T_1, T_2\}$ there exists a set $E' \subseteq E(N)$ of reticulation edges such that T can be derived from N by conducting the following steps.

(1) Delete each edge from N that is not contained in E'.

(2) Remove each node whose corresponding taxon is not contained in $\mathcal{X}'$.
(3) Remove each unlabeled node of out-degree 0 repeatedly.
(4) Suppress each node of both in- and out-degree 1.

Notice that, for determine E', it suffices to select at most one in-edge of each hybridization node in N.

From a biological point of view, this means that N displays T if each speciation event of T is reflected by N. Moreover, each internal node of in-degree 1 represents a speciation event and each internal node providing an in-degree of at least 2 represents a reticulation event or, in terms of hybridization, a hybridization event. This means, in particular, that such a latter node represents an individual whose genome is a *chimaera* of several parents. Thus, such a node v of in-degree larger than or equal to 2 is called a *hybridization node* and each edge directed into v is called a *hybridization edge*.

Now, based on those hybridization nodes, the *reticulation number* $r(N)$ of a hybridization network N is defined by

$$r(N) = \sum_{v \in V : \delta^-(v) > 0} \big(\delta^-(v) - 1\big) = |E| - |V| + 1, \tag{1}$$

where V denotes the node set and E the edge set of N. Next, based on the definition of the reticulation number, for two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2 the hybridization number $h(\{T_1, T_2\})$ is

$$\min\{r(N) : \text{N is a hybridization network displaying } T_1 \text{ and } T_2\}. \tag{2}$$

Lastly, we call a *hybridization network* N for two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2 a *minimum hybridization network* if $r(N) = h(\{T_1, T_2\})$.

Forests. Let T be a rooted nonbinary phylogenetic $\mathcal{X}$-tree T. Then, we call any set of rooted nonbinary phylogenetic trees $\mathcal{F} = \{F_1, \ldots, F_k\}$ with $\mathcal{L}(\mathcal{F}) = \mathcal{X}$ a *forest on* $\mathcal{X}$, if we have for each pair of trees F_i and F_j that $\mathcal{L}(F_i) \cap \mathcal{L}(F_j) = \emptyset$. Moreover, if additionally for each component F in $\mathcal{F}$ the tree $T|_{\mathcal{L}(F)}$ equals F, we say that $\mathcal{F}$ is a *forest for* T. Lastly, let $\mathcal{F}$ be a forest for a rooted binary phylogenetic $\mathcal{X}$-tree T. Then, by $\overline{\mathcal{F}}$ we refer to the forest that is obtained from $\mathcal{F}$ by deleting each component only consisting of an isolated node as well as the element containing the node labeled by taxon ρ if it contains at most one edge.

Agreement Forests. For technical purpose, the definition of agreement forests is based on two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2 whose roots are marked by a unique taxon $\rho \notin \mathcal{X}$ as follows. Let r_i be the root of the tree T_i with $i \in \{1, 2\}$. Then, we first create a new node v_i labeled by a new taxon $\rho \notin \mathcal{X}$ and then attach this node to r_i by inserting the edge (v_i, r_i). Notice that this case v_1 and v_2 is the new root of T_1 and T_2, respectively. Moreover, since we consider ρ as being a new taxon, the taxa set of both trees is $\mathcal{X} \cup \{\rho\}$.

Now, given two such marked rooted binary phylogenetic trees T_1 and T_2 on $\mathcal{X} \cup \rho$, a set of components $\mathcal{F} = \{F_\rho, F_1, \ldots, F_k\}$ is an *agreement forest* for T_1 and T_2 if the following three conditions are satisfied.

(1) Each component F_i with taxa set $\mathcal{X}_i$ equals $T_1|_{\mathcal{X}_i}$ and $T_2|_{\mathcal{X}_i}$.
(2) There is exactly one component, denoted as F_ρ, with $\rho \in \mathcal{L}(F_\rho)$.
(3) Let $\mathcal{X}_\rho, \mathcal{X}_1, \ldots, \mathcal{X}_k$ be the respective taxa sets of $F_\rho, F_1, \ldots, F_k$. All trees in $\{T_1(\mathcal{X}_i) | i \in \{\rho, 1, \ldots, k\}\}$ as well as $\{T_2(\mathcal{X}_i) | i \in \{\rho, 1, \ldots, k\}\}$ are node disjoint subtrees of T_1 and T_2, respectively.

An illustration of an agreement forest is given in Figure 1(b).

Lastly, an agreement forest for two rooted binary phylogenetic $\mathcal{X}$-trees containing a minimum number of components is called a *maximum agreement forest*. Now, based on maximum agreement forests, one can compute a minimum hybridization network if an additional constraint is satisfied, which is presented in the following. Notice that, from a biological point of view, this constraint is necessary since it prevents species from inheriting genetic material from their own offspring.

Modified ancestor descendant graphs. Given two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2 together with a forest $\mathcal{F} = \{F_\rho, F_1, \ldots, F_k\}$ for T_1 (or T_2), the modified ancestor descendant graph $AG^*(T_1, T_2, \mathcal{F})$ with node set $\mathcal{F}$ contains a directed edge (F_i, F_j), $i \neq j$, if

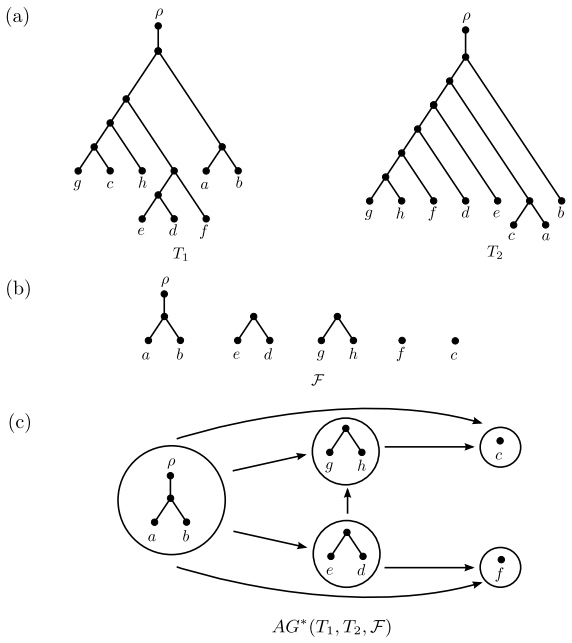

Fig. 1. (a) Two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2 with taxa set $\mathcal{X} = \{a, b, c, d, e, f, g, h, \rho\}$. (b) An acyclic agreement forest $\mathcal{F}$ for T_1 and T_2. (c) The directed graph $AG^*(T_1, T_2, \mathcal{F})$ not containing any directed cycles and, thus, $\mathcal{F}$ is acyclic.

(1) the root of $T_1(\mathcal{L}(F_i))$ is an ancestor of the root of $T_1(\mathcal{L}(F_j))$,
(2) or the root of $T_2(\mathcal{L}(F_i))$ is an ancestor of the root of $T_2(\mathcal{L}(F_j))$.

An illustration of such an modified ancestor descendant graph is given in Figure 1(c).

Now, we say that $\mathcal{F}$ is an *acyclic agreement forest* if $\mathcal{F}$ satisfies the conditions of an agreement forest and if the corresponding graph $AG^*(T_1, T_2, \mathcal{F})$ does not contain any *directed* cycles. Again, each acyclic agreement forest of minimum size among all acyclic agreement forests is called a *maximum acyclic agreement forest*. Note that the definition of an ancestor descendant graph given in both works Baroni *et. al.* [3] as well as Scornavacca *et. al.* [12] is different to the one presented above. Regarding the definitions of those two works, the set of components $\mathcal{F}$ has to be an agreement forest for T_1 and T_2 and, thus, this definition is more strict than the one given for $AG^*(T_1, T_2, \mathcal{F})$ in which $\mathcal{F}$ just has to be a forest for T_1 (or T_2). For a better illustration, in Figure 1 we give an example of an acyclic agreement forest together with its underlying modified ancestor descendant graph.

Notice that the concept of maximum acyclic agreement forests for two rooted binary phylogenetic $\mathcal{X}$-trees is of interest, since its size minus one corresponds to the hybridization number of these two trees, which is stated by Theorem 1 given in the paper of Baroni *et. al.* [3].

Theorem 1 ([3]). *Let $\mathcal{F}$ be a maximum acyclic agreement forest of size k for two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2, then $h(T_1, T_2) = k$.*

Cherries. For a rooted binary phylogenetic $\mathcal{X}$-tree T we call two of its leaves a and c a *cherry*, denoted by $\{\mathcal{L}(a), \mathcal{L}(c)\}$, if both nodes have the same parent. Moreover, let R and S be two rooted binary phylogenetic trees on $\mathcal{X}_R$ and $\mathcal{X}_S$, respectively, so that $\mathcal{X}_R \subseteq \mathcal{X}_S$, and let $\mathcal{F}$ be a forest for S. Then, a cherry $\{\mathcal{L}(a), \mathcal{L}(c)\}$ in R is called a *common cherry of R and $\mathcal{F}$*, if there exists a cherry $\{\mathcal{L}(a'), \mathcal{L}(c')\}$ in $\mathcal{F}$ with $\mathcal{L}(a') = \mathcal{L}(a)$ and $\mathcal{L}(c') = \mathcal{L}(c)$. Otherwise, the cherry $\{\mathcal{L}(a), \mathcal{L}(c)\}$ in R is called *contradicting cherry of R and $\mathcal{F}$*. Moreover, in order to ease reading, if there exists a component in $\mathcal{F}$ containing two leaves labeled by $\mathcal{L}(a)$ and $\mathcal{L}(c)$, respectively, we write $a \sim_{\mathcal{F}} c$. Otherwise, if such a component does not exist, we write $a \not\sim_{\mathcal{F}} c$.

Moreover, let $\mathcal{F}$ be a forest for a rooted binary phylogenetic $\mathcal{X}$-tree and let $\{\mathcal{L}(a), \mathcal{L}(c)\}$ be a cherry that is contained in $\mathcal{F}$ in which e denotes the in-edge of its parent p. Then, if e exists, by $\mathcal{F} \div \{\mathcal{L}(a), \mathcal{L}(c)\}$ we simply refer to the forest $\mathcal{F} - \{e\}$ and, otherwise, if p has in-degree 0, to $\mathcal{F}$. Furthermore, let $\mathcal{P}$ be the path connecting a and c in $\mathcal{F}$. Then, the *set of pendant edges for* $\{\mathcal{L}(a), \mathcal{L}(c)\}$ contains each edge (v, w) with $v \in V(\mathcal{P}) \setminus \{a, c\}$ and $w \notin V(\mathcal{P})$, where $V(\mathcal{P})$ denotes the set of nodes in $\mathcal{P}$. Notice that in general there exist several, namely precisely $|V(\mathcal{P})| - 3$, edges satisfying the condition of such an edge.

Cherry Reductions. Let $\mathcal{F}$ be a forest for a rooted binary phylogenetic $\mathcal{X}$-tree and let $\{\mathcal{L}(a), \mathcal{L}(c)\}$ be a cherry of a component F_i in $\mathcal{F}$. Then, a *cherry reduction*, according to a cherry $\{\mathcal{L}(a), \mathcal{L}(c)\}$ in one of its components F_i, implies the following two operations.

(1) The parent of the two nodes a and c is labeled by $\{\mathcal{L}(a), \mathcal{L}(c)\}$.
(2) Both nodes a and c together with their adjacent edges are deleted from F_i.

Throughout the algorithm, such a reduction step of a common cherry $\{\mathcal{L}(a), \mathcal{L}(c)\}$ in $\mathcal{F}$ is shortly denoted by $\mathcal{F}[\{\mathcal{L}(a), \mathcal{L}(c)\} \to \mathcal{L}(a) \cup \mathcal{L}(c)]$. Furthermore, the reverse notation $\mathcal{F}[\mathcal{L}(a) \cup \mathcal{L}(c) \to \{\mathcal{L}(a), \mathcal{L}(c)\}]$ describes the insertion of two new taxa labeled by $\mathcal{L}(a)$ and $\mathcal{L}(c)$, respectively, together with removing the label $\mathcal{L}(a) \cup \mathcal{L}(c)$ from its parent. Notice that, for applying such insertion steps, one has to keep track of the preceding reduction steps.

Equivalently, for a cherry $\{\mathcal{L}(a), \mathcal{L}(c)\}$ in a rooted binary phylogenetic $\mathcal{X}$-tree R we write $R[\{\mathcal{L}(a), \mathcal{L}(c)\} \to \mathcal{L}(a) \cup \mathcal{L}(c)]$ to denote a cherry reduction of $\{\mathcal{L}(a), \mathcal{L}(c)\}$ in R.

3 The Algorithm allMAAFs

In this section, we present the algorithm allMAAFs that was first published in the work of Scornavacca *et al.* Scornavacca *et. al.* [12]. To increase its readability, we decided to split the original algorithm allMAAFs into six parts (cf. Alg. 1–5). Note that, apart from its graphical representation, our presentation of the algorithm allMAAFs together with its terminology in general adheres to the original algorithm.

Algorithm 1: allMAAFs(S, T, R, $\mathcal{F}$, k, M)

Data: Two rooted binary phylogenetic $\mathcal{X}$-trees S and T, a rooted binary phylogenetic tree R and a forest $\mathcal{F}$ such that $\mathcal{L}(R) = \mathcal{L}(\mathcal{F})$ and $\mathcal{L}(T) = \mathcal{L}(\mathcal{F})$, an integer k, and a list M that contains information of previously reduced cherries.
Result: A set $\boldsymbol{\mathcal{F}}$ of forests for $\mathcal{F}$ and an integer. In particular, if $\mathcal{F} = T$, $R = S$, $M = \emptyset$, and $k \geq h(S, T)$ is the input to allMAAFs, the output precisely consists of all maximum-acyclic-agreement forests for S and T and their respective hybridization number.

```
1  if k < 0 then
2      return (∅, k − 1);
3  if |L(R)| = 0 then
4      F' ← cherryExpansion(F, M);
5      if AG(S, T, F') is acyclic then
6          return (F', |F'| − 1);
7      else
8          return (∅, k − 1);
9  else
10     let {L(a), L(c)} be a cherry of R;
11     if {L(a), L(c)} is a common cherry of R and F then
12         return (ProcessCommonCherry(S, T, R, F, k, M, {L(a), L(c)}));
13     if k ≠ (|F| − 1) or {L(a), L(c)} is a contradicting cherry of R and F then
14         return (ProcessContradictingCherry(S, T, R, F, k, M, {L(a), L(c)}));
```

Algorithm 2: cherryExpansion($\mathcal{F}, M$)

while M is not empty **do**
  $M \leftarrow$ remove last element of M, say$\{\mathcal{L}(a), \mathcal{L}(c)\}$;
  $\mathcal{F} \leftarrow \mathcal{F}[\mathcal{L}(a) \cup \mathcal{L}(c) \rightarrow \{\mathcal{L}(a), \mathcal{L}(c)\}]$;
return $\mathcal{F}$

Algorithm 3: cherryReduction($R, \mathcal{F}, M, \{\mathcal{L}(a), \mathcal{L}(c)\}$)

$M' \leftarrow$ Add $\{\mathcal{L}(a), \mathcal{L}(c)\}$ as last element of M;
$R' \leftarrow R[\{\mathcal{L}(a), \mathcal{L}(c)\} \rightarrow \mathcal{L}(a) \cup \mathcal{L}(c)]$;
$\mathcal{F}' \leftarrow \mathcal{F}[\{\mathcal{L}(a), \mathcal{L}(c)\} \rightarrow \mathcal{L}(a) \cup \mathcal{L}(c)]$;
return $(R', \mathcal{F}', M')$

Algorithm 4: ProcessCommonCherry(S, T, R, $\mathcal{F}$, k, M, $\{\mathcal{L}(a), \mathcal{L}(c)\}$)

$(R', \mathcal{F}', M') \leftarrow$ cherryReduction(R, $\mathcal{F}$, M, $\{\mathcal{L}(a), \mathcal{L}(c)\}$);
$(\boldsymbol{\mathcal{F}_r}, k_r) \leftarrow$ allMAAFs(S, T, $R'|_{\mathcal{L}(\overline{\mathcal{F}'})}$, $\mathcal{F}'$, k, M');
if $\boldsymbol{\mathcal{F}_r} \neq \emptyset$ **then**
  $k \leftarrow \min(k, k_r)$;
if $(k = |\mathcal{F}| - 1)$ **then return** $(\boldsymbol{\mathcal{F}_r}, k)$;
else
  $(\boldsymbol{\mathcal{F}_a}, k_a) \leftarrow$ allMAAFs(S, T, $R|_{\mathcal{L}(\overline{\mathcal{F}-\{e_a\}})}$, $\mathcal{F} - \{e_a\}$, $k-1$, M);
  if $\boldsymbol{\mathcal{F}_a} \neq \emptyset$ **then**
    $k \leftarrow \min(k, k_a - 1)$;
  if $(k_a - 1 = k)$ **then** $\boldsymbol{\mathcal{F}} \leftarrow \boldsymbol{\mathcal{F}} \cup \boldsymbol{\mathcal{F}_a}$;
  $(\boldsymbol{\mathcal{F}_c}, k_c) \leftarrow$ allMAAFs(S, T, $R|_{\mathcal{L}(\overline{\mathcal{F}-\{e_c\}})}$, $\mathcal{F} - \{e_c\}$, $k-1$, M);
  if $\boldsymbol{\mathcal{F}_c} \neq \emptyset$ **then**
    $k \leftarrow \min(k, k_c - 1)$;
  $\boldsymbol{\mathcal{F}} \leftarrow \emptyset$;
  if $(k_a - 1 = k)$ **then** $\boldsymbol{\mathcal{F}} \leftarrow \boldsymbol{\mathcal{F}_a}$;
  if $(k_c - 1 = k)$ **then** $\boldsymbol{\mathcal{F}} \leftarrow \boldsymbol{\mathcal{F}} \cup \boldsymbol{\mathcal{F}_c}$;
  if $(k_r = k)$ **then** $\boldsymbol{\mathcal{F}} \leftarrow \boldsymbol{\mathcal{F}} \cup \boldsymbol{\mathcal{F}_r}$;
  return $(\boldsymbol{\mathcal{F}}, k)$;

Algorithm 5: ProcessContradictingCherry(S, T, R, $\mathcal{F}$, k, M, $\{\mathcal{L}(a), \mathcal{L}(c)\}$)

$(\boldsymbol{\mathcal{F}_a}, k_a) \leftarrow$ allMAAFs(S, T, $R|_{\mathcal{L}(\overline{\mathcal{F}-\{e_a\}})}$, $\mathcal{F} - \{e_a\}$, $k-1$, M);
if $\boldsymbol{\mathcal{F}_a} \neq \emptyset$ **then**
  $k \leftarrow \min(k, k_a - 1)$;
$(\boldsymbol{\mathcal{F}_c}, k_c) \leftarrow$ allMAAFs(S, T, $R|_{\mathcal{L}(\overline{\mathcal{F}-\{e_c\}})}$, $\mathcal{F} - \{e_c\}$, $k-1$, M);
if $\boldsymbol{\mathcal{F}_c} \neq \emptyset$ **then**
  $k \leftarrow \min(k, k_c - 1)$;
$\boldsymbol{\mathcal{F}} \leftarrow \emptyset$;
if $a \nsim_{\mathcal{F}} c$ **then**
  if $(k_a - 1 = k)$ **then** $\boldsymbol{\mathcal{F}} \leftarrow \boldsymbol{\mathcal{F}_a}$;
  if $(k_c - 1 = k)$ **then** $\boldsymbol{\mathcal{F}} \leftarrow \boldsymbol{\mathcal{F}} \cup \boldsymbol{\mathcal{F}_c}$;
  return $(\boldsymbol{\mathcal{F}}, k)$;
else
  $(\boldsymbol{\mathcal{F}_B}, k_B) \leftarrow$ allMAAFs(S, T, $R|_{\mathcal{L}(\overline{\mathcal{F}-\{e_B\}})}$, $\mathcal{F} - \{e_B\}$, $k-1$, M);
  if $\boldsymbol{\mathcal{F}_B} \neq \emptyset$ **then**
    $k \leftarrow \min(k, k_B - 1)$;
  if $(k_a - 1 = k)$ **then** $\boldsymbol{\mathcal{F}} \leftarrow \boldsymbol{\mathcal{F}_a}$;
  if $(k_B - 1 = k)$ **then** $\boldsymbol{\mathcal{F}} \leftarrow \boldsymbol{\mathcal{F}} \cup \boldsymbol{\mathcal{F}_B}$;
  if $(k_c - 1 = k)$ **then** $\boldsymbol{\mathcal{F}} \leftarrow \boldsymbol{\mathcal{F}} \cup \boldsymbol{\mathcal{F}_c}$;
  return $(\boldsymbol{\mathcal{F}}, k)$;

4 Modifications to allMAAFs

In this section, we present two modifications of the algorithm allMAAFs that, on the one hand, do not improve its theoretical runtime but, one the other hand, significantly improve its practical running time, which will be indicated by a simulation study reported in Section 6.

The first modification improves the processing of a certain type of contradicting cherry whereas for the other two modifications only the processing step of a common cherry is of interest.

Given a contradicting cherry $\{\mathcal{L}(a), \mathcal{L}(c)\}$ for R and $\mathcal{F}$ with $a \sim_{\mathcal{F}} c$, the original algorithm conducts three recursive calls. One by recursively calling the algorithm with $\mathcal{F}-\{e_a\}$ and $R|_{\overline{\mathcal{F}-\{e_a\}}}$, one by recursively calling the algorithm with $\mathcal{F}-\{e_c\}$ and $R|_{\overline{\mathcal{F}-\{e_c\}}}$ and one by recursively calling the algorithm with $\mathcal{F}-\{e_B\}$ and $R|_{\overline{\mathcal{F}-\{e_B\}}}$, in which e_a and e_c refers to the in-edge of leaf a and c, respectively, in $\mathcal{F}$ and e_B refers to an in-edge of a subtree lying on the path connecting a and c in $\mathcal{F}$. Regarding the latter recursive call, in the upcoming part of this work we will show that, in order to compute all maximum acyclic agreement forests, instead of cutting just one in-edge e_B one can cut all of those in-edges all at once.

Moreover, given a common cherry $\{\mathcal{L}(a), \mathcal{L}(c)\}$ for R and $\mathcal{F}$, the original algorithm allMAAFs always branches into three new computational paths; one path corresponding to the cherry reduction of $\{\mathcal{L}(a), \mathcal{L}(c)\}$ and two corresponding to the deletion of both in-edges of the two leaves a and c (cf. Alg. 4, line 7–15). To understand the sense of our two modifications, one has to take the necessity of these two additional edge deletions into account.

Therefor, we demonstrate a specific scenario that is outlined in Figure 2 showing two phylogenetic trees T_1 and T_2 as well as a maximum acyclic agreement $\mathcal{F} = \{((a, b), \rho), (e, d), (g, h), c, f\}$ for those two trees. By running the algorithm allMAAFs for T_1 and T_2 without deleting in-edges of a common cherry, the given maximum acyclic agreement forest $\mathcal{F}$ is never computed. This is due to the fact that, once the component $F_i = ((g, h), f)$ occurs on any computational path, F_i will be part of the resulting agreement forest. Because of F_i and the component (e, d), such a resulting agreement forest is never acyclic and, thus, does not satisfy the conditions of an acyclic agreement forest. However, by cutting instead of contracting the common cherry $\{(g, h), f\}$, the resulting agreement forest turns into the maximum acyclic agreement forest $\mathcal{F}$. This example implies that sometimes the deletion of in-edges corresponding to taxa of a common cherry is necessary, which is, however, in practice not often the case, and, thus, the original algorithm allMAAFs usually produces a lot of additional unnecessary computational steps compared with our second modification offering a different solution for such a scenario.

4.1 The Algorithm allMAAFs[1]

Our first algorithm allMAAFs[1] is a modification of the original algorithm allMAAFs improving the processing of contradicting cherries. Let $\{\mathcal{L}(a), \mathcal{L}(c)\}$ be a contradicting cherry of R and $\mathcal{F}$ such that $a \sim_{\mathcal{F}} c$ holds and let e_B be an edge that is defined as follows. Let $\mathcal{P}$ be the path connecting a and c in $\mathcal{F}$. Then, the edge set E_B contains each edge $e_B = (v, w)$ with $v \in V(\mathcal{P}) \setminus \{a, c\}$ and $w \notin V(\mathcal{P})$, where $V(\mathcal{P})$ denotes the set of nodes of $\mathcal{P}$. Note that in this case there exist precisely $|V(\mathcal{P})-3|$ edges satisfying the conditions of such an edge e_B. Now, if such a cherry $\{\mathcal{L}(a), \mathcal{L}(c)\}$ occurs, the original algorithm allMAAFs branches into a computational path by cutting exactly one of those edges in E_B. We will show, however, that in this case the whole set E_B can be cut from $\mathcal{F}$ all at once without having an impact on the computation of all maximum acyclic agreement forests for both input trees. In Algorithm 6, we give a pseudo code of allMAAFs[1]. Note that, for the sake of clarity, we just present the modified part in respect of the original algorithm dealing with the processing of contradicting cherries. The remaining parts are unmodified and can be looked up in Section 3.

4.2 The Algorithm allMAAFs[2]

Our third algorithm allMAAFs[2] is again a modification of our first algorithm allMAAFs[1] and is based on a tool turning agreement forests into acyclic agreement forests. This tool, published by Whidden *et al.* Whidden *et. al.* [14], is based on the concept of an *expanded cycle graph* refining cyclic agreement forest. Due to such

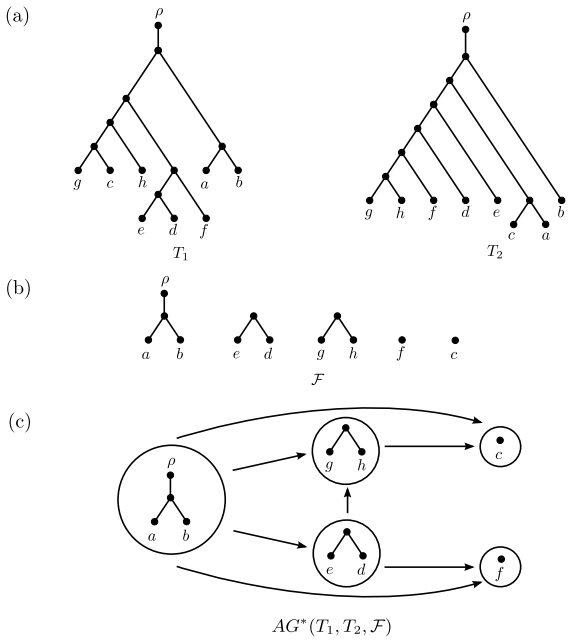

Fig. 2. (a) Two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2 with taxa set $\mathcal{X} = \{a, b, c, d, e, f, g, h, \rho\}$. (b) An acyclic agreement forest $\mathcal{F}$ for T_1 and T_2. (c) The directed graph $AG^*(T_1, T_2, \mathcal{F})$ not containing any directed cycles and, thus, $\mathcal{F}$ is acyclic.

Algorithm 6: ProcessContradictingCherry$^1(S,\ T,\ R,\ \mathcal{F},\ k,\ M,\ \{\mathcal{L}(a), \mathcal{L}(c)\})$

```
(F_a, k_a) ← allMAAFs¹(S, T, R|_{L(F−{e_a})}, F − {e_a}, k − 1, M);
if F_a ≠ ∅ then
    k ← min(k, k_a − 1);
(F_c, k_c) ← allMAAFs¹(S, T, R|_{L(F−{e_c})}, F − {e_c}, k − 1, M);
if F_c ≠ ∅ then
    k ← min(k, k_c − 1);
F ← ∅;
if a ≁_F c then
    if (k_a − 1 = k) then F ← F_a;
    if (k_c − 1 = k) then F ← F ∪ F_c;
    return (F, k);
else
    (F_B, k_B) ← allMAAFs¹(S, T, R|_{L(F−E_B)}, F − E_B, k − |E_B|, M);
    if F_B ≠ ∅ then
        k ← min(k, k_B − 1);
    if (k_a − 1 = k) then F ← F_a;
    if (k_B − 1 = k) then F ← F ∪ F_B;
    if (k_c − 1 = k) then F ← F ∪ F_c;
    return (F, k);
```

additional refinement steps, which are performed right after the computation of each maximum agreement forest, both cutting steps for processing a common cherry can be omitted. The simulation study in Section 6 indicates that this refinement step is efficient enough so that this modification in general outperforms the original algorithm allMAAFs as well as our first modification allMAAFs^1. In Algorithm 7 and 8, we present a pseudo code describing the algorithm allMAAFs^2. Again, for the sake of clarity, we restrict the presentation to only those parts that are modified in respect to allMAAFs^1. Moreover, we omit a description of the subroutine conducting the refinement step, denoted by RefineForest, as it can be looked up in the work of Whidden *et al.* Whidden *et. al.* [14]. Notice that, due to this refinement step, an implementation of this modification is quite more expensive compared with our first modified algorithm allMAAFs^1.

Algorithm 7: allMAAFs$^2(S,\, T,\, R,\, \mathcal{F},\, k,\, M)$

```
1  if k < 0 then
2      return (∅, k − 1);
3  if |L(R)| = 0 then
4      F' ← cherryExpansion(F, M);
5      F'' ← RefineForest(F');
6      F ← ∅;
7      foreach F'' ∈ F'' do
8          if |F''| = k then
9              F ← F ∪ F'';
10         else if |F''| < k then
11             k ← |F''|;
12             F ← {F''};
13     return (F, k − 1);
14 else
15     let {L(a), L(c)} be a cherry of R;
16     if {L(a), L(c)} is a common cherry of R and F then
17         return (ProcessCommonCherry^3(S, T, R, F, k, M, {L(a), L(c)}));
18     if k ≠ (|F| − 1) or {L(a), L(c)} is a contradicting cherry of R and F then
19         return (ProcessContradictingCherry(S, T, R, F, k, M, {L(a), L(c)}));
```

Algorithm 8: ProcessCommonCherry$^3(S,\, T,\, R,\, \mathcal{F},\, k,\, M,\, \{\mathcal{L}(a), \mathcal{L}(c)\})$

```
1  (R', F', M') ← cherryReduction(R, F, M, {L(a), L(c)});
2  (F_r, k_r) ← allMAAFs^2(S, T, R'|_{L(F')‾}, F', k, M');
3  if (k_r = k) then return (F_r, k);
4  return (∅, k);
```

5 Proofs of Correctness

In this section, we give formal proofs showing the correctness of all two modified algorithms presented in Section 4. In a first step, however, we give some further definitions that are crucial for what follows.

5.1 The algorithm processCherries

In the following, we introduce the algorithm processCherries, which has already been utilized in the paper of Scornavacca *et al.* Scornavacca *et. al.* [12]. This algorithm is a simplified version of the algorithm allMAAFs describing one of its computational paths by a list of *cherry actions*. Notice that this algorithm is a major tool that will help us to establish the correctness of our two modified algorithms.

Cherry actions. Let R be a rooted binary phylogenetic $\mathcal{X}$-tree and let $\mathcal{F}$ be a forest for R. Then, a cherry action $\wedge = (\{\mathcal{L}(a), \mathcal{L}(c)\}, e)$ is a tuple containing a set $\{\mathcal{L}(a), \mathcal{L}(c)\}$ of two taxa of two leaves a and c as well as an edge e. We say that $\wedge$ is a *cherry action for R and $\mathcal{F}$*, if $\{\mathcal{L}(a), \mathcal{L}(c)\}$ is a cherry of R and if, additionally, one of the following three conditions is satisfied.

(1) Either $\{\mathcal{L}(a), \mathcal{L}(c)\}$ is a common cherry of R and $\mathcal{F}$ and $e \in \{\emptyset, e_a, e_c\}$,
(2) or $\{\mathcal{L}(a), \mathcal{L}(c)\}$ is a contradicting cherry of R and $\mathcal{F}$ with $a \not\sim_{\mathcal{F}} c$ and $e \in \{e_a, e_c\}$,
(3) or $\{\mathcal{L}(a), \mathcal{L}(c)\}$ is a contradicting cherry of R and $\mathcal{F}$ with $a \sim_{\mathcal{F}} c$ and $e \in \{e_a, e_B, e_c\}$.

In this context, e_a and e_c is an edge in $\mathcal{F}$ adjacent to both leaves a and c, respectively. Moreover, e_B is part of the set of pendant edges for $\{\mathcal{L}(a), \mathcal{L}(c)\}$.

Cherry lists. Now, given two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2, we say that $\bigwedge$ is a *cherry list for T_1 and T_2*, if, while calling processCherries$(T_1, \{T_2\}, \bigwedge)$, in the i-th iteration $\wedge_i$ is a cherry action for R_i and $\mathcal{F}_i$. Note that, if $\bigwedge$ is not a cherry list for T_1 and T_2, calling processCherries$(T_1, \{T_2\}, \bigwedge)$ (cf. Alg. 9) returns the empty set.

Algorithm 9: processCherries$(R, \mathcal{F}, \bigwedge = (\wedge_1, \ldots, \wedge_n))$

```
1  M ← ∅;
2  foreach i ∈ 1, ..., n do
3      if ∧_i is a cherry action for R and F then
4          ({L(a), L(c)}, e_i) ← ∧_i;
5          if e_i = ∅ then
6              M ← Add {L(a), L(c)} as last element of M;
7              R ← R[{L(a), L(c)} → L(a) ∪ L(c)];
8              F ← F[{L(a), L(c)} → L(a) ∪ L(c)];
9          else
10             F ← F − {e_i};
11             R ← R|_{L(F)};
12     else
13         return ∅;
14 while M is not empty do
15     M ← remove the last element, say {L(a), L(c)}, from M;
16     F ← F[L(a) ∪ L(c) → {L(a), L(c)}];
17 return R, F, M
```

Further details as well as an example of processCherries can be found in the work of Scornavacca *et al.* Scornavacca *et. al.* [12].

Next, we will introduce the algorithm processCherries[1]. This algorithm is a simplified version of the algorithm allMAAFs[1] describing one of its computational paths by a list of *extended cherry actions*. Notice

that this algorithm is a major tool that will help us to establish the correctness of our first modified algorithm allMAAFs^1.

Extended cherry action. Let R be a rooted binary phylogenetic $\mathcal{X}$-tree and let $\mathcal{F}$ be a forest for R. Then, an extended cherry action $\wedge = (\{\mathcal{L}(a), \mathcal{L}(c)\}, E')$ is a tuple containing a set $\{\mathcal{L}(a), \mathcal{L}(c)\}$ of two taxa of two leaves a and c as well as an edge set E'. We say that $\wedge$ is an extended cherry action for R and $\mathcal{F}$, if $\{\mathcal{L}(a), \mathcal{L}(c)\}$ is a contradicting cherry of R and $\mathcal{F}$ with $a \sim_{\mathcal{F}} c$ and $E' = E_B$, where E_B refers to the set of pendant edges for $\{\mathcal{L}(a), \mathcal{L}(c)\}$ in $\mathcal{F}$.

Extended cherry list. Now, given two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2, we say that $\bigwedge$ is an *extended cherry list for T_1 and T_2*, if, while calling processCherries$^1(T_1, \{T_2\}, \bigwedge)$, in the i-th iteration $\wedge_i$ is either a cherry action or an extended cherry action for R_i and $\mathcal{F}_i$. Note that, if $\bigwedge$ is not an extended cherry list for T_1 and T_2, calling processCherries$^1(T_1, \{T_2\}, \bigwedge)$ (cf. Alg. 10) returns the empty set.

Algorithm 10: processCherries$^1(R, \mathcal{F}, \bigwedge = (\wedge_1, \ldots, \wedge_n))$

```
M ← ∅;
foreach i ∈ 1, ..., n do
    if ∧_i is a cherry action for R and F then
        ({L(a), L(c)}, e_i) ← ∧_i;
        if e_i = ∅ then
            M ← Add {L(a), L(c)} as last element of M;
            R ← R[{L(a), L(c)} → L(a) ∪ L(c)];
            F ← F[{L(a), L(c)} → L(a) ∪ L(c)];
        else
            F ← F − {e_i};
            R ← R|_{L(F)};
    else if ∧_i is an extended cherry action for R and F then
        ({L(a), L(c)}, E_i) ← ∧_i;
        F ← F − E_i;
        R ← R|_{L(F)};
    else
        return ∅;
while M is not empty do
    M ← remove the last element, say {L(a), L(c)}, from M;
    F ← F[L(a) ∪ L(c) → {L(a), L(c)}];
return R, F, M
```

Lemma 1. *Let $\bigwedge'$ be an extended cherry list for two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2. Moreover, let $\mathcal{F}$ be a forest for T_2 calculated by $\bigwedge'$. Then, there also exists a cherry list $\bigwedge$ for T_1 and T_2 calculating $\mathcal{F}$.*

Proof. To proof this lemma we will show how to replace each extended cherry action of $\bigwedge'$ so that the resulting cherry list $\bigwedge$ still computes $\mathcal{F}$. Let $\wedge'_i = (\{\mathcal{L}(a), \mathcal{L}(c)\}, E_B)$ be an extended cherry action with $E_B = \{e_1, e_2, \ldots, e_k\}$. Then, we can replace $\wedge'_i$ through the sequence of cherry actions

$$(\{\mathcal{L}(a), \mathcal{L}(c)\}, e_1), (\{\mathcal{L}(a), \mathcal{L}(c)\}, e_2), \ldots, (\{\mathcal{L}(a), \mathcal{L}(c)\}, e_k).$$

Since by these cherry actions the same edges are cut from T_2 as by $\wedge'_i$, the topology of R'_{i+1} equals R_{i+k+1}, which directly implies that $\mathcal{F}$ is still computed.

5.2 Correctness of allMAAFs^1

In this section, we will discuss the correctness of our first modified algorithm allMAAFs^1 by establishing the following theorem.

Theorem 2. *Let T_1 and T_2 be two rooted binary phylogenetic $\mathcal{X}$-trees and $k \in \mathbb{N}$. Calling*

$$\text{allMAAFs}^1(T_1, T_2, T_1, \{T_2\}, k, \emptyset)$$

returns all maximum acyclic agreement forests for T_1 and T_2 if and only if $k \geq h(T_1, T_2)$.

Proof. Let T_1 and T_2 be two rooted binary phylogenetic $\mathcal{X}$-trees and let $\bigwedge$ be a cherry list for T_1 and T_2 mimicking a computational path of allMAAFs calculating a maximum acyclic agreement forest $\mathcal{F}$ for T_1 and T_2. Then, in the following, we say a cherry action $\wedge_i = (\{\mathcal{L}(a), \mathcal{L}(c)\}, e_i)$ in $\bigwedge$ is a *special cherry action*, if $\{\mathcal{L}(a), \mathcal{L}(c)\}$ is a contradicting cherry of R_i and $\mathcal{F}_i$ and if e_i is contained in the set of pendant edges for $\{\mathcal{L}(a), \mathcal{L}(c)\}$. Note that, whereas such special cherry actions may occur in a computational path of allMAAFs, this is not the case for a computational path of allMAAFs^1. In the following, however, we will show that the algorithm allMAAFs calculates a maximum acyclic agreement forest $\mathcal{F}$ for T_1 and T_2 if and only if $\mathcal{F}$ is calculated by allMAAFs^1.

Lemma 2. *Let T_1 and T_2 be two rooted binary phylogenetic $\mathcal{X}$-trees and let $\mathcal{F}$ be a maximum acyclic agreement forest for T_1 and T_2 of size k. Then, $\mathcal{F}$ is calculated by calling* allMAAFs$(T_1, T_2, T_1, T_2, k, \emptyset)$ *if and only if $\mathcal{F}$ is calculated by calling* allMAAFs$^1(T_1, T_2, T_1, \{T_2\}, k, \emptyset)$.

Proof. '$\Longrightarrow$': From Lemma 1 we can directly deduce that if there exists a computational path of allMAAFs^1 calculating $\mathcal{F}$, then, there also exists a computational path in allMAAFs calculating $\mathcal{F}$.

'$\Longleftarrow$': Let $\wedge_j = (\{\mathcal{L}(a), \mathcal{L}(c)\}, e_j = \emptyset)$ be a cherry action of $\bigwedge$ contracting both leaves a and c. Then, we say a preceding special cherry action $\wedge_i = (\{\mathcal{L}(x), \mathcal{L}(y)\}, e_i)$ $(i < j)$ refers to $\wedge_k$ if the following condition is satisfied. Let a_i and c_i be the lowest common ancestor of $\mathcal{L}(a)$ and $\mathcal{L}(c)$ in R_i, respectively. Then, $e_i \neq \emptyset$ has to be a pendant edge lying on the path connecting both nodes a_i and c_i. In such a case, we say the special cherry action is either of *Type A* or *Type B* (see definition below). Otherwise, if a special cherry action does not refer to another cherry action, we say this cherry action is of *Type C*.

Now, based on the edge $e_i = (v, w)$, being part of the special action $\wedge_i$ referring to the cherry action $\wedge_j$, we further distinguish whether $\wedge_i$ is either of *Type A* or of *Type B*. Therefore, let $T(w)$ be the subtree rooted at w and let $\mathcal{L}(T(w))$ be the set of taxa being contained in $T(w)$. Now, we say $\wedge_i$ is of *Type A* if there exists a forest $\mathcal{F}_k$ and a tree R_k with $i < k < j$ so that both of the following two conditions are satisfied.

(i) Each taxon in $\mathcal{L}(T(w))$ is part of a taxa set of an isolated node. Notice that, as a direct consequence, there exists a leaf w' in R_k with label $\mathcal{L}(w)$.
(ii) The sibling s' of w' in R_k is a leaf. Notice that, as a direct consequence, $\{\mathcal{L}(s'), \mathcal{L}(w')\}$ is a cherry of R_k.

Otherwise, if there does not exist such a forest $\mathcal{F}_k$ and such a tree R_k satisfying these two conditions, we say $\wedge_i$ is of *Type B*. An illustration of these two types is given in Figure 3

Notice that, due to the so chosen definition, each special cherry action in $\bigwedge$ has to be either of *Type A*, of *Type B*, or of *Type C*. Next, based on this observation, we will show in three steps how to turn $\bigwedge$ into another cherry list for T_1 and T_2 not containing any special cherry actions, but still calculating $\mathcal{F}$, which can be briefly summarized as follows.

In a first step, we will show how to modify $\bigwedge$ by replacing each special cherry action of *Type A* and of *Type C* through a non-special cherry action so that the result is still a cherry list for T_1 and T_2 calculating $\mathcal{F}$. Next, during a second and a third step, we will show how to replace each set of special cherry actions of *Type B* all referring to the same cherry action by a single extended cherry action so that the result is an extended cherry list for T_1 and T_2 still calculating $\mathcal{F}$.

Step 1. Let $\wedge_i = (\{\mathcal{L}(x), \mathcal{L}(y)\}, e_i = (v, w))$ be a special cherry action of *Type A* or of *Type C*, and let $\wedge_k$ with $i < k < j$ be the first cherry action of $\bigwedge$ in which in $\mathcal{F}_k$ each taxon in $\mathcal{L}(T(w))$ is contained in a taxa set of an isolated node so that in R_k the sibling s' of the leaf labeled by $\mathcal{L}(w)$ is also leaf. Then, let $\bigwedge'$

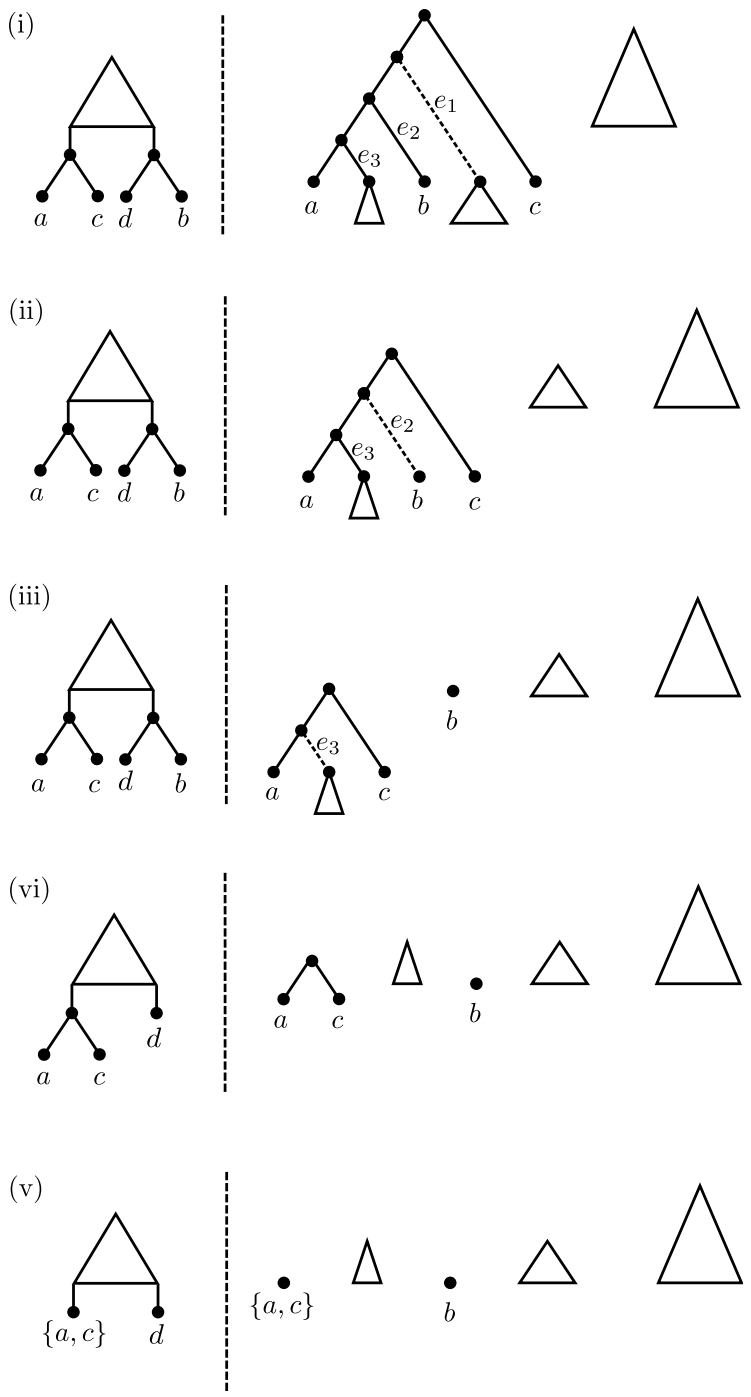

Fig. 3. An illustration of the cherry list $((\{a,c\},e_1),(\{a,c\}),e_2),(\{a,c\}),e_3),(\{a,c\}),\emptyset))$ in which $(\{a,c\},e_1)$ is a special cherry action of *Type B* and $(\{a,c\},e_2)$ is a special cherry action of *Type A*. Note that the expanded cherry list $((\{b,d\},e_2),(\{a,c\},\{e_1,e_3\}),(\{a,c\}),\emptyset))$ applied to (i) yields the same scenario as depicted in (v) without making use of any special cherry actions.

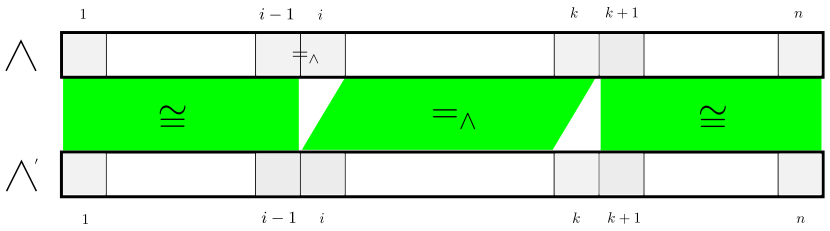

Fig. 4. An illustration of the scenario as described in Step 1. The figure shows a comparison between two sequences of R_i-trees; one corresponding to the original cherry list and one corresponding to the modified cherry list, in which a special cherry action of *Type A* is shifted from position i to $k+1$.

be a cherry list that is obtained from $\bigwedge$ by first removing $\wedge_i = (\{\mathcal{L}(x), \mathcal{L}(y)\}, e_i)$ and then by inserting the cherry action $(\{\mathcal{L}(w'), \mathcal{L}(s')\}, e_{w'})$ right after $\wedge_k$, so that $\bigwedge'$ equals

$$(\wedge'_1, \ldots, \wedge'_n) = (\wedge_1, \ldots, \wedge_{i-1}, \wedge_{i+1}, \ldots, \wedge_k, \wedge'_k, \ldots, \wedge'_n),$$

where $\wedge'_k = (\{\mathcal{L}(w'), \mathcal{L}(s')\}, e_{w'})$ with $e_{w'}$ being the in-edge of w'.

Now, for the following, remember that R_l (or R'_l) refers to the input tree occurring during iteration l while processing the cherry list $\bigwedge$ (or $\bigwedge'$). Moreover, we write $R_i =_\wedge R'_j$ if both trees contain the same set of cherries. Then, based on the position of the cherry action $\wedge_l$ in $\bigwedge$, we can make the following five observations (cf. Fig. 4).

- If $l < i$, then R'_l equals R_l: This is the case because R_1 equals R'_1 and through $\wedge_l$ and $\wedge'_l$ the same tree operation is performed on R_l and R'_l, respectively.
- If $l = i$, then $R'_{l-1} =_\wedge R_l$: This is the case because R_{i-1} equals R'_{i-1}, in R_l the node w cannot be part of a cherry (due to the definition of a special cherry action of *Type A*) and, thus, through $\wedge_l$ $(=\wedge_i)$ the set of cherries in R_{i-1} remains unchanged.
- If $i < l < k+1$, then $R'_{l-1} =_\wedge R_l$: This is the case because $R_i =_\wedge R'_{i-1}$ and again in R_l node w cannot be part of a cherry. Thus, the two cherry actions $\wedge_l$ and $\wedge'_{l-1}$ produce the same set of cherries in R_l and R'_{l-1}.
- If $l = k+1$, then R'_{k+1} equals R_{k+1}: This is the case because $R_k =_\wedge R'_{k-1}$ and through $\wedge'_k$ first node w is cut and then removed from R'_{k-1}. Consequently, through the cherry actions $\wedge_q$ and $\wedge'_q$, with $1 \leq q < k+1$, the the same tree operations are applied to R_1 and R'_1 and, thus, R_{k+1} equals R'_{k+1}.
- If $l > k+1$, then R'_l equals R_l: This is the case because R_{k+1} equals R'_{k+1} and through $\wedge_l$ and $\wedge'_l$ the same tree operation is performed on R_l and R'_l, respectively.

Due to these observations and since through $\bigwedge'$ still the same edges are cut from T_2 as through the original cherry list $\bigwedge$, $\bigwedge'$ is a cherry list for T_1 and T_2 still calculating $\mathcal{F}$. Now, by consecutively replacing all special cherry actions of *Type A* and *Type C* we can turn $\bigwedge$ into the cherry list $\bigwedge^{(1)}$ for T_1 and T_2 only containing special cherry actions of *Type B* and still computing $\mathcal{F}$. Next, we will show how to remove each of those remaining special cherry actions.

Step 2. Let $\bigwedge^* \subset \bigwedge$ with $|\bigwedge^*| = k$ be a set of special cherry actions of *Type B* all referring to a cherry action $\wedge_j = (\{\mathcal{L}(a), \mathcal{L}(c)\}, \emptyset)$, and let $i = \min_{i'}\{\wedge_{i'} : \wedge_{i'} \in \bigwedge^*\}$ with $\wedge_i = (\{\mathcal{L}(x), \mathcal{L}(z)\}, e_i)$. Moreover, let $\bigwedge'$ be the cherry list that is obtained from $\bigwedge^{(1)}$ as follows. First the cherry of each cherry action in $\bigwedge^*$ is replaced by $\{\mathcal{L}(a), \mathcal{L}(c)\}$ and then all those cherry actions in $\bigwedge^*$ are rearranged such that they are placed in sequential order directly right before position j.

This means, in particular, that $\bigwedge'$ contains a sequence of special cherry actions

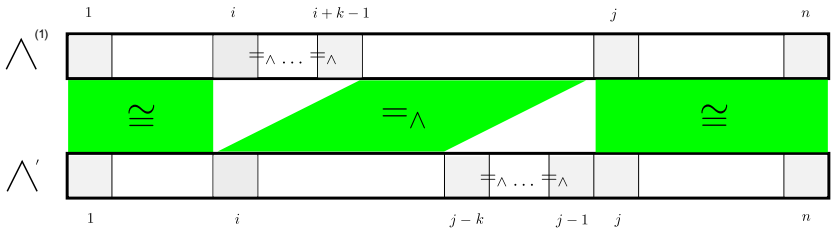

Fig. 5. An illustration of the scenario as described in Step 2. The figure shows a comparison of the two sequences of R_i-trees; one corresponding to the original cherry list and one corresponding to the modified cherry list, in which a sequence of k special cherry actions of *Type B*, beginning at position i, is shifted to position $j-k$.

$$
\begin{aligned}
&\vdots\\
\wedge_{j-k} &= (\{\mathcal{L}(a), \mathcal{L}(c)\}, e_{j-k} = (v_{j-k}, w_{j-k})),\\
\wedge_{j-k+1} &= (\{\mathcal{L}(a), \mathcal{L}(c)\}, e_{j-k+1} = (v_{j-k+1}, w_{j-k+1})),\\
&\vdots\\
\wedge_{j-1} &= (\{\mathcal{L}(a), \mathcal{L}(c)\}, e_{j-1} = (v_{j-1}, w_{j-1})),\\
\wedge_{j} &= (\{\mathcal{L}(a), \mathcal{L}(c)\}, \emptyset),\\
&\vdots
\end{aligned}
$$

in which each edge of $\bigwedge^*$ is contained in the set of pendant edges for the cherry $\{\mathcal{L}(a), \mathcal{L}(c)\}$. Now, just for convenience, in the following we assume, without loss of generality, that all those cherry actions in $\bigwedge^*$ occur in sequential order beginning at position i. Moreover, for the following, let $\mathcal{W}$ be the set of target nodes of each edge in $\bigwedge^*$ and let R_l (or R'_l) be the input tree of iteration l while processing the cherry list $\bigwedge^{(1)}$ (or $\bigwedge'$). Additionally, again we write $R_i =_\wedge R'_j$ if both trees contain the same set of cherries. Then, based on the position of a cherry action $\wedge_l$ in $\bigwedge$, we can make the following five observations (cf. Fig. 5).

- If $l < i$, then R'_l equals R_l: This is the case because R_1 equals R'_1 and through $\wedge_l$ and $\wedge'_l$ the same tree operation is performed on R_l and R'_l, respectively.
- If $i-1 < l < i+k$, then $R_l =_\wedge R'_{i-1}$: This is the case because R_{i-1} equals R'_{i-1} and, since in R_l each w_i in $\mathcal{W}$ cannot be part of a cherry, through $\wedge_l$ the set of cherries in R_l remains unchanged.
- If $i+k-1 < l < j$, then $R'_{l-k} =_\wedge R_l$: This is the case because $R_{i+k-1} =_\wedge R'_{i-1}$ and again in R_l each w_i in $\mathcal{W}$ cannot be part of a cherry. Thus, the two cherry actions $\wedge_l$ and $\wedge'_{l-k}$ produce the same set of cherries in R_l and R'_{l-k}.
- If $l = j$, then R'_l equals R_l: This is the case because $R_{j-1} =_\wedge R'_{j-k-1}$ and through each cherry action in $\bigwedge^*$ each subtree $T(w_i)$ is cut (and removed from R'_{j-k-1}, if w_i is a leaf). Consequently, through the cherry actions $\wedge_q$ and $\wedge'_q$, with $1 \leq q < j$, the same tree operations are applied to R_1 and R'_1 and, thus, R_j equals R'_j.
- If $l > j$, then R'_l equals R_l: This is the case because R_j equals R'_j and through $\wedge_l$ and $\wedge'_l$ the same tree operation is performed on R_l and R'_l, respectively.

Now, by consecutively rearranging all special cherry actions of *Type B* as described above, we can turn $\bigwedge$ into the cherry list $\bigwedge^{(2)}$ for T_1 and T_2 in which all special cherry actions referring to the same cherry action are located next to each other. Moreover, as a direct consequence of these observations and since through $\bigwedge^{(2)}$ still the same edges are cut from T_2 as through $\bigwedge^{(1)}$, $\bigwedge^{(2)}$ is still a cherry list for T_1 and T_2 calculating $\mathcal{F}$.

Step 3. Let $\bigwedge^{(2)}$ be the list that is obtained from applying Step 1 and Step 2 as described above. Then, we can further modify the cherry list $\bigwedge^{(2)}$ to $\bigwedge^{(3)}$ by replacing each sequence

$$(\wedge_i = (\{\mathcal{L}(a), \mathcal{L}(c)\}, e_i), \ldots, \wedge_{i+k-1} = (\{\mathcal{L}(a), \mathcal{L}(c)\}, e_{i+k-1}))$$

of special cherry actions (i.e., cherry actions of *Type B*) through a single extended cherry action $\wedge_B = (\{\mathcal{L}(a), \mathcal{L}(c)\}, E_B)$, where E_B denotes the pendant edge set for $\{\mathcal{L}(a), \mathcal{L}(c)\}$. As a consequence, since $\wedge_B$ just summarizes all tree operations conducted by the replaced sequence of special cherry actions, by $\bigwedge^{(3)}$ the same edges are cut from T_2 as it is the case for the cherry list $\bigwedge^{(2)}$. As a direct consequence, the maximum acyclic agreement forest $\mathcal{F}$ is still calculated by the extended cherry list $\bigwedge^{(3)}$.

In summary, as described by those three steps, we can consecutively replace all special cherry actions of $\bigwedge$ so that the resulting list $\bigwedge^{(3)}$ satisfies all of the following three conditions.

- $\bigwedge^{(3)}$ is an extended cherry list for T_1 and T_2.
- $\bigwedge^{(3)}$ does not contain any special cherry actions.
- $\bigwedge^{(3)}$ calculates the maximum acyclic agreement forest $\mathcal{F}$.

This directly implies that for each computational path in allMAAFs calculating a maximum acyclic agreement forest $\mathcal{F}$ for both input trees there also exists a computational path in allMAAFs^1 calculating $\mathcal{F}$.

Now, from Lemma 2 we can directly deduce the correctness of Theorem 2.

5.3 Correctness of allMAAFs^2

The correctness of our third modified algorithm allMAAFs^2 principally directly follows from the correctness of the refinement step, which can be found in the work of Whidden *et al.* Whidden *et. al.* [14]. We still have to show, however, that omitting both additional cutting steps when processing a common cherry still computes all of those agreement forests from which all maximum acyclic agreement forest can be obtained by cutting some of its edges.

Theorem 3. *Let T_1 and T_2 be two rooted binary phylogenetic $\mathcal{X}$-trees and $k \in \mathbb{N}$. Calling*

$$\text{allMAAFs}^2(T_1, T_2, T_1, \{T_2\}, k, \emptyset)$$

returns all maximum acyclic agreement forests for T_1 and T_2 if and only if $k \geq h(T_1, T_2)$.

Proof. To proof Theorem 3 we first have to establish a further lemma. Here, we argument that by contracting common cherries instead of cutting one of its edges, results in an agreement forests containing less components.

Lemma 3. *Let $\bigwedge = (\wedge_1, \wedge_2, \ldots, \wedge_n)$ be a cherry list for two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2 with $\wedge_i = (\{\mathcal{L}(a), \mathcal{L}(c)\}, e_c) \in \bigwedge$ so that $\{\mathcal{L}(a), \mathcal{L}(c)\}$ is a common cherry of R_i and $\mathcal{F}_i$ and e_c the in-edge of leaf c. Moreover, let $\mathcal{F}$ be the agreement forest for T_1 and T_2 that is calculated by calling* processCherries$(R_1, \mathcal{F}_1, \bigwedge)$ *such that $|\mathcal{F}| = k$. Then, there exists an agreement forest $\hat{\mathcal{F}}$ with $|\hat{\mathcal{F}}| = k - 1$ that can be computed by calling* processCherries$(R_i, \mathcal{F}_i, (\wedge'_i, \ldots, \wedge'_n))$ *in which $\wedge'_i = (\{\mathcal{L}(a), \mathcal{L}(c)\}, \emptyset)$.*

Proof. Let F_a and F_c be the two components in $\mathcal{F}$ that have been derived from expanding both components containing a and c, respectively, and let $\mathcal{X}_a$ be the taxa set of the subtree rooted at a (i.e., $F_a(a)$). Then, by attaching F_c back to F_a we can reduce the size of the agreement forest $\mathcal{F}$ by one. Here, depending on whether $\text{LCA}_{F_a}(\mathcal{X}_a)$ is the root of F_a or not, this re-attachment step can be done in two different ways.

– If $\mathrm{LCA}_{F_a}(\mathcal{X}_a)$ is not the root of F_a and, thus, has an in-going edge $e = (u, w)$, first e is split into two adjacent edges (u, v) and (v, w) and then F_c is re-attached to v by inserting a new edge (v, r_c), where r_c denotes the root of F_c.
– Otherwise, if $\mathrm{LCA}_{F_a}(\mathcal{X}_a)$ is the root of F_a, first a new node v is created and then v is connected to the two roots of F_a and F_c.

In the following, we will denote the component that is obtained from attaching F_c back to F_a by F_r and the resulting set of components by $\hat{\mathcal{F}}$. In Figure 6, we give an illustration of the two forests $\mathcal{F}$ and $\hat{\mathcal{F}}$.

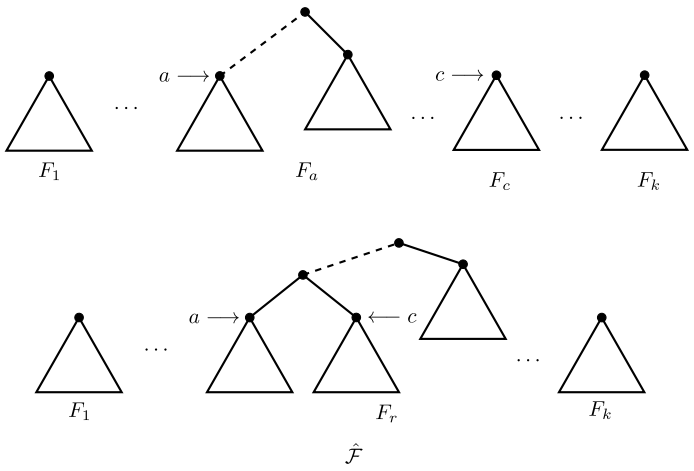

Fig. 6. The two forests $\mathcal{F}$ and $\hat{\mathcal{F}}$ as defined in Lemma 3. Note that, since F_c is attached to F_a, the size of $\hat{\mathcal{F}}$ is $|\mathcal{F}| - 1$.

Notice that we can calculate the agreement forest $\hat{\mathcal{F}}$ by calling the algorithm processCherries with a specific cherry list $\bigwedge'$ that can be derived from the original cherry list $\bigwedge = (\wedge_1, \wedge_2, \ldots, \wedge_i = (\{\mathcal{L}(a), \mathcal{L}(c)\}, e_c), \ldots, \wedge_n)$ as follows.

(1) Replace $\wedge_i$ by $(\{\mathcal{L}(a), \mathcal{L}(c)\}, \emptyset)$.
(2) In each subsequent cherry action $\wedge_{i+1}, \ldots, \wedge_n$, replace $\mathcal{L}(a)$ through $\mathcal{L}(a) \cup \mathcal{L}(c)$.

Let R_{i+1} and $\mathcal{F}_{i+1}$ be computed by applying $(\{\mathcal{L}(a), \mathcal{L}(c)\}, e_c)$ to R_i and $\mathcal{F}_i$ and let R'_{i+1} and $\mathcal{F}'_{i+1}$ be computed by applying $(\{\mathcal{L}(a), \mathcal{L}(c)\}, \emptyset)$ to R_i and $\mathcal{F}_i$. Since $\{\mathcal{L}(a), \mathcal{L}(c)\}$ is a common cherry of R_i and $\mathcal{F}_i$, the only difference between the two topologies, disregarding node labels, is that $\mathcal{F}_{i+1}$ contains an additional component consisting of an isolated node labeled by $\mathcal{L}(c)$. Notice, however, that, since this component is fully contracted, this component cannot have an impact on the subsequent cherry actions $\wedge_{i+1}, \ldots, \wedge_n$. Consequently, the cherry list $\bigwedge'$ calculates $\hat{\mathcal{F}}$.

Up to now, we have shown that there exists a set of components $\hat{\mathcal{F}}$ of size $|\mathcal{F}|-1$, which can be computed by calling the algorithm processCherries with a slightly modified cherry list $\bigwedge'$. Now, to establish Lemma 3, we still have to show that $\hat{\mathcal{F}}$ is an agreement forest for both input trees T_1 and T_2. First, note that each expanded forest to which allMAAFs applies the acyclic check (cf. Alg. 1, line 5) has to satisfy each condition of an agreement forest [12, Lemma 3]. As $\bigwedge'$ is imitating a computational path corresponding to the original algorithm allMAAFs calculating $\hat{F}$, this directly implies that $\hat{F}$ is an agreement forest for T_1 and T_2.

Now, given two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2, assume that our first modified algorithm allMAAFs[1] contains a computational path calculating a maximum acyclic agreement forest $\mathcal{F}$ of size k by cutting instead of contracting a common cherry. More precisely, let $\bigwedge$ with $\wedge_i = (\{\mathcal{L}(a), \mathcal{L}(c)\} \in \bigwedge$ be the cherry list mimicking this computational path and, without loss of generality, let $\wedge_i$ be the only cherry

action cutting instead of contracting a common cherry. Then, as already discussed in the proof of Lemma 3, there additionally exists a cherry list $\bigwedge'$ calculating a specific agreement forest $\hat{\mathcal{F}}$ of size $k-1$ containing a component $\hat{F}$. More specifically, let F_a and F_c be the expanded component in $\mathcal{F}$ which is derived from the contracted node a and c, respectively. Then, $\hat{\mathcal{F}}$ is obtained from $\mathcal{F}$ by re-attaching F_c back to F_a (cf. Fig. 6). Thus, as a direct consequence, by cutting the in-edge corresponding to the root of $\hat{F}(\mathcal{L}(F_c))$ the maximum acyclic agreement forest $\mathcal{F}$ arises.

This implies that, if cutting instead of contracting a common cherry yields a maximum acyclic agreement forest $\mathcal{F}$, the algorithm allMAAFs^2 guarantees the computation of an agreement forest $\hat{\mathcal{F}}$ from which $\mathcal{F}$ can be obtained by cutting some of its edges. Note that the refinement step is always able to identify the minimal number of such edges and, thus, as the only difference between the first modified algorithm and the algorithm allMAAFs^2 consists in the way a common cherry is processed, Theorem 3 is established.

6 Simulation Study

Our simulation study has been conducted on the same synthetic dataset as the one used for the simulation study reported in the work of Albrecht *et al.* Albrecht *et. al.* [1]. It consists of binary phylogenetic tree pairs that have been generated in respect to one combination of the following three parameters: the *number of taxa* $n = \{20, 50, 100.200\}$, the *executed number of rSPR-moves* $k = \{5, 10, \ldots, 50\}$, and the *tangling degree* $d = \{3, 5, 10, 15, 20\}$. For each combination of those parameters 10 tree pairs have been generated, resulting in 2000 tree pairs in total.

More precisely, a tree pair (T_1, T_2) is computed as follows. In a first step, the first tree T_1 with n leaves is computed which is done initially by randomly selecting two nodes u and v of a specific set $\mathcal{V}$ consisting of n nodes of both in- and out-degree 0. Then, those two selected nodes u and v are connected to a new node w and, finally, $\mathcal{V}$ is updated by replacing u and v by its parent w. This process is repeated until $\mathcal{V}$ consists only of one node corresponding to the root of T_1. In a subsequent step, T_2 is computed by applying k rSPR-moves to T_1 each respecting tangling degree d.

When performing a sequence of rSPR-moves, one can undo or redo some of those moves and, thus, k is only an upper bound of the real underlying rSPR-distance corresponding to both trees of a tree pair. The tangling degree, as already described in the work of Albrecht *et al.* Albrecht *et. al.* [1], is an *ad hoc* concept controlling the number of minimum common clusters during the construction of a tree pair. Since, however, all the four implementations that are tested on the synthetic dataset perform a cluster reduction and, thus, equally benefit from the number of minimum common clusters, this number is irrelevant for our simulation study and, consequently, we do not give any further details about this parameter. Instead, we refer the interested reader to the work of Albrecht *et al.* Albrecht *et. al.* [1].

In this section, the practical runtimes produced by the respective implementations of our two modifications are compared with the practical runtimes of an implementation of the original algorithm. More precisely, by applying our synthetic dataset the practical runtime was measured for an implementation of the original algorithm allMAAFs, the algorithm allMAAFs^1, and the algorithm allMAAFs^2. Each of those algorithms has been integrated as a plug-in into the freely available Java software Hybroscale[1]. The simulation study has been conducted on a grid computer providing 16 cores and 40 GB RAM. In order to receive a reasonable set of completed data sets within an appropriate time period, we decided to compute only the hybridization number and omitted the computation of all maximum acyclic agreement forests. Moreover, we set the maximum runtime of each tree pair to 20 minutes, which means that each tree pair whose computation of the hybridization number could not be finished within 20 minutes was aborted. Notice that depending on the runtime analysis, which is shown in the following, those aborted tree pairs were either not taken into account or counted as being finished after 20 minutes.

The problem arising when computing all, instead of just one, maximum acyclic agreement forests for certain tree sets is that, typically, there exists a large number of those agreement forests all being distributed in a vast search space. Consequently, even if those tree sets are of low computational complexity, one has to

[1] `www.bio.ifi.lmu.de/softwareservices/hybroscale`

investigate far more than 20 minutes, which means that one could not conduct the simulation study within an appropriate time period. Otherwise, if we would choose a relatively small maximum running time, such as 20 minutes, we could only process those tree sets of low computational complexity, which are not able to indicate the speedup obtained from applying our three modifications.

In Figure 7, the mean average runtimes in terms of the computed hybridization numbers are shown. More precisely, this plot was generated by first aggregating all tree pairs corresponding to the same hybridization number and then by computing the mean average runtime of each of those aggregated subsets. Therefore, each aborted tree pair whose computation would have taken longer than the given time limit of 20 minutes, was treated as follows. If the tree pair could not be computed by both considered algorithms, this tree pair was not taken into account. However, if at least one algorithm was able to compute the hybridization number for a tree pair, the runtime of the other algorithm was set to 20 minutes if this algorithm was aborted in this case. Regarding the Figure 7, the number assigned to each measurement denotes the number of tree pairs whose hybridization number could be computed within the time limit by the corresponding two algorithm.

Figure 7 indicates that each of our two modified algorithms outperforms the original algorithm. Moreover, this figure indicates that our third modified algorithm allMAAFs^2 is significantly more efficient than our first modified algorithm allMAAFs^1.

Figure 8 shows the distributions of the speedups obtained from our two modified algorithms via boxplots. Each of those boxplots was generated by first selecting a relevant set $\mathcal{D}$ of tree pairs from our synthetic dataset and then by computing the speedup of each of those tree pairs by taking the runtime obtained from the corresponding two algorithms into account. More precisely, in a first step, we set the runtime of each tree pair whose hybridization number could not be computed within the time limit to 20 minutes. Second, we consider each tree pair d as relevant if at least one of both algorithms could process d within the time limit of 20 minutes and if at least one computation referring to one of both algorithms took longer than 50 seconds. Consequently, the relevant set $\mathcal{D}$ excludes those tree pairs whose computational complexity is, on the one hand, too low and, on the other hand, too high to reveal a difference between the runtimes of both algorithms.

Figure 8 again indicates that both modified algorithms are more efficient than the original algorithm. More specifically, for the considered set of tree pairs, allMAAFs^1 and allMAAFs^2 is on mean average about 3.9 times and 11.9 times (median 2.3, 4.9 and 6.6), respectively, faster than allMAAFs. Moreover, our second modified algorithm can significantly improve the practical runtime of our first modified algorithm. More specifically, for the considered set of tree pairs, allMAAFs^2 is on mean average 8.15 times (median 3.18), respectively, faster than allMAAFs^1.

In order to give a reason of the speedup obtained by our two modified algorithms, we measured the number of recursive calls that have been performed for calculating the hybridization number of each tree pair. To draw comparisons, we only took those tree pairs into account whose hybridization number could be processed by both algorithms within the time limit of 20 minutes, which were in total 1302. Next, all these tree pairs were grouped according to their hybridization number and, finally, the mean average number of recursive calls for each group was computed. Regarding Figure 9, the numbers that are attached to each measurement correspond to the number of tree pairs that have contributed to the mean average denoted at the x-axis.

Figure 9 indicates that by applying our two modified algorithms there are significantly less recursive calls necessary for the computation of a maximum acyclic agreement forest compared with the original algorithm. Moreover, Figure 9 (c) shows that our second modified algorithm allMAAFs^2 has to conduct significantly less recursive calls than our first modified algorithm for computing hybridization numbers which, obviously, compensates the effort of preventing additional recursive calls when processing a common cherry. Since the computation of the hybridization number is just an intermediate step in computing all maximum acyclic agreement forests, the difference between the number of recursive calls between the first and both the second and the third modified algorithm is expected to be even larger in this case.

Finally, we finish this section by comparing the runtimes of the four algorithms via scatter-plots (cf. Fig. 10). Again, all tree pairs whose hybridization number could not be computed within the time limit was set to

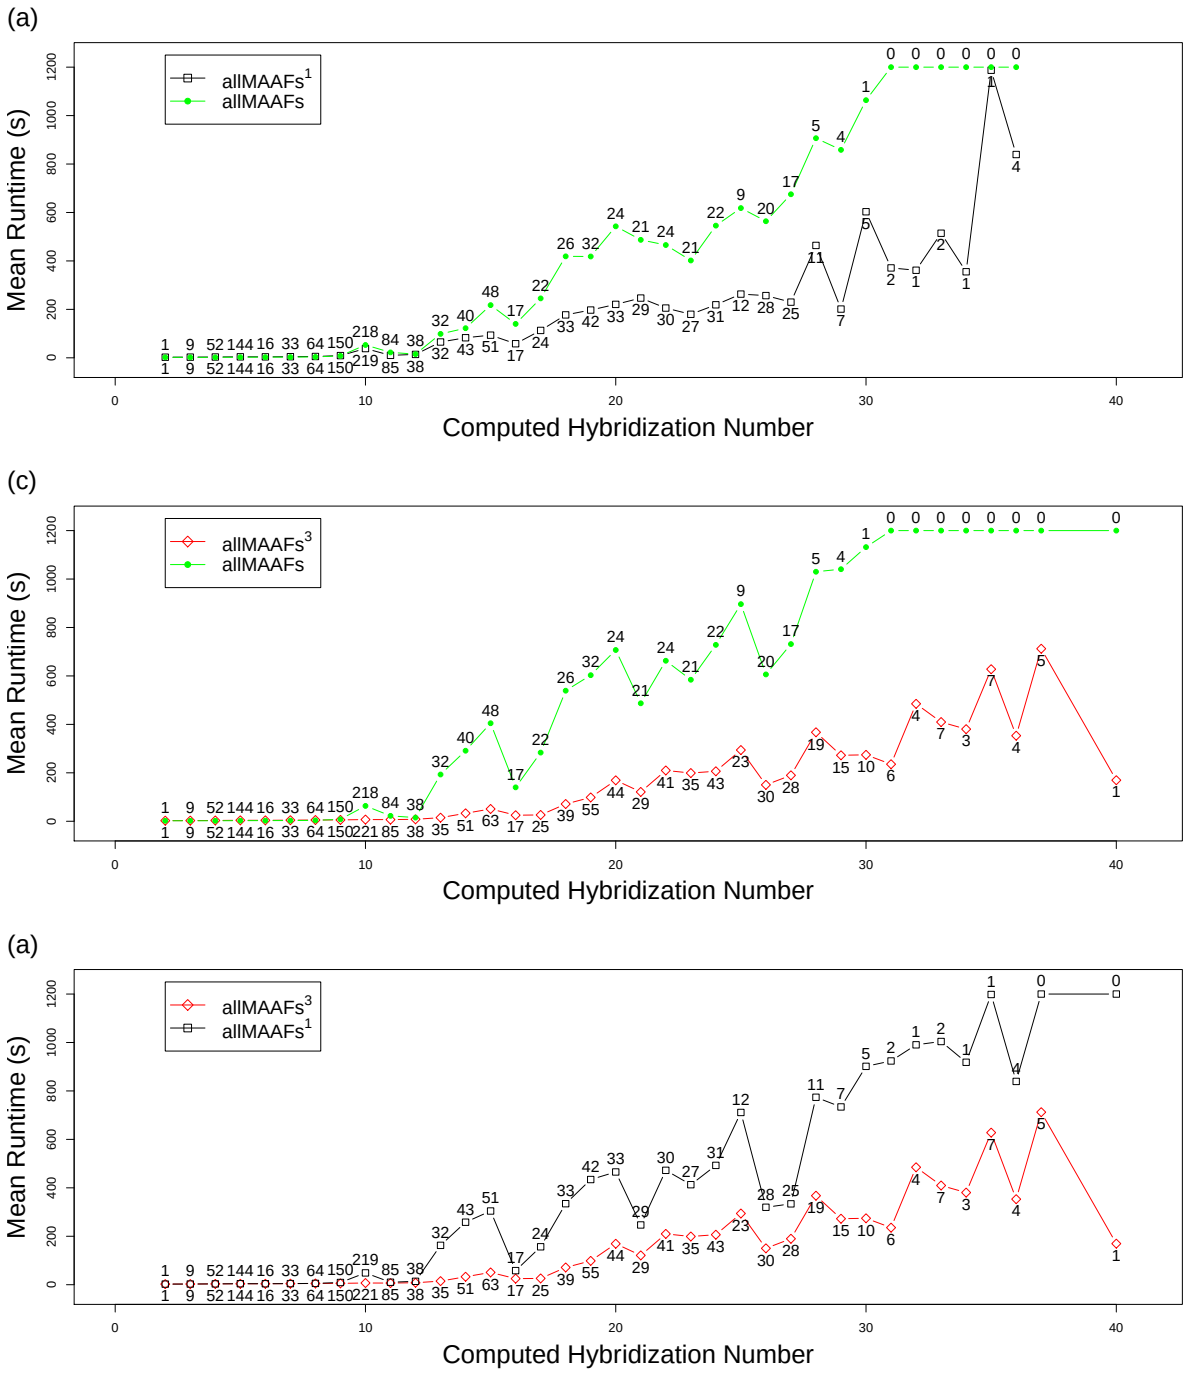

Fig. 7. Comparisons between the mean average runtimes in terms of the hybridization number of (a) allMAAFs and allMAAFs[1],(b) allMAAFs and allMAAFs[2], as well as (c) allMAAFs[1] and allMAAFs[2]. For each hybridization number the corresponding number of tree pairs is given that could be computed within the time limit of 20 minutes by the two respective algorithms and, thus, contributed to the mean average runtime denoted by the y-axis.

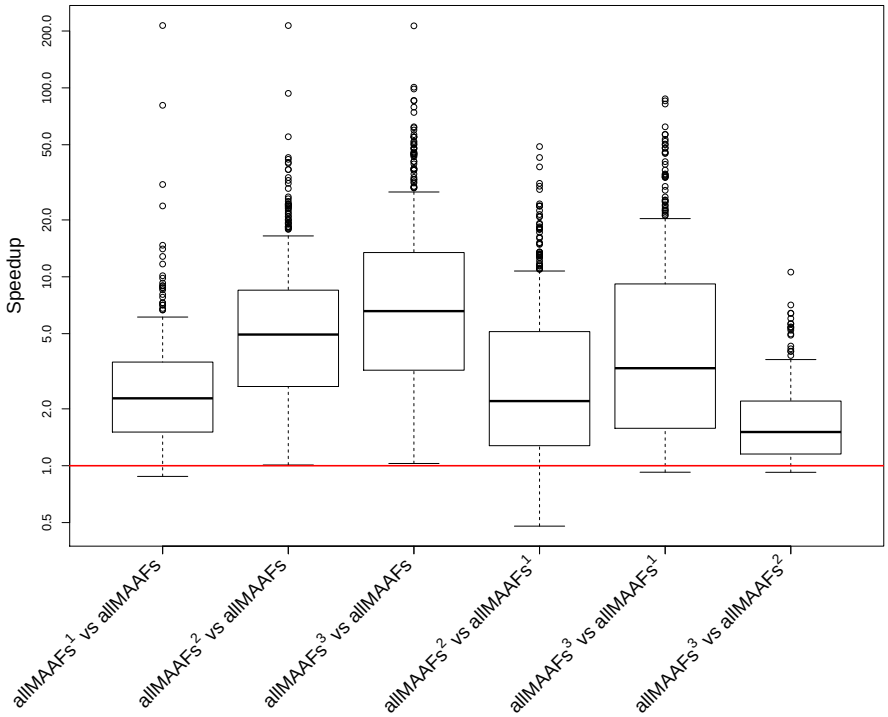

Fig. 8. The distributions of the speedups, obtained from comparing algorithm A_1 versus algorithm A_2, as denoted at the x-axis, via boxplots. A tree pair d from our synthetic dataset was only considered if at least one of the two algorithms could process d within the time limit of 20 minutes and if at least one algorithm took longer than 50 seconds for its computation. Notice that a base-10 log scale is used for the y-axis.

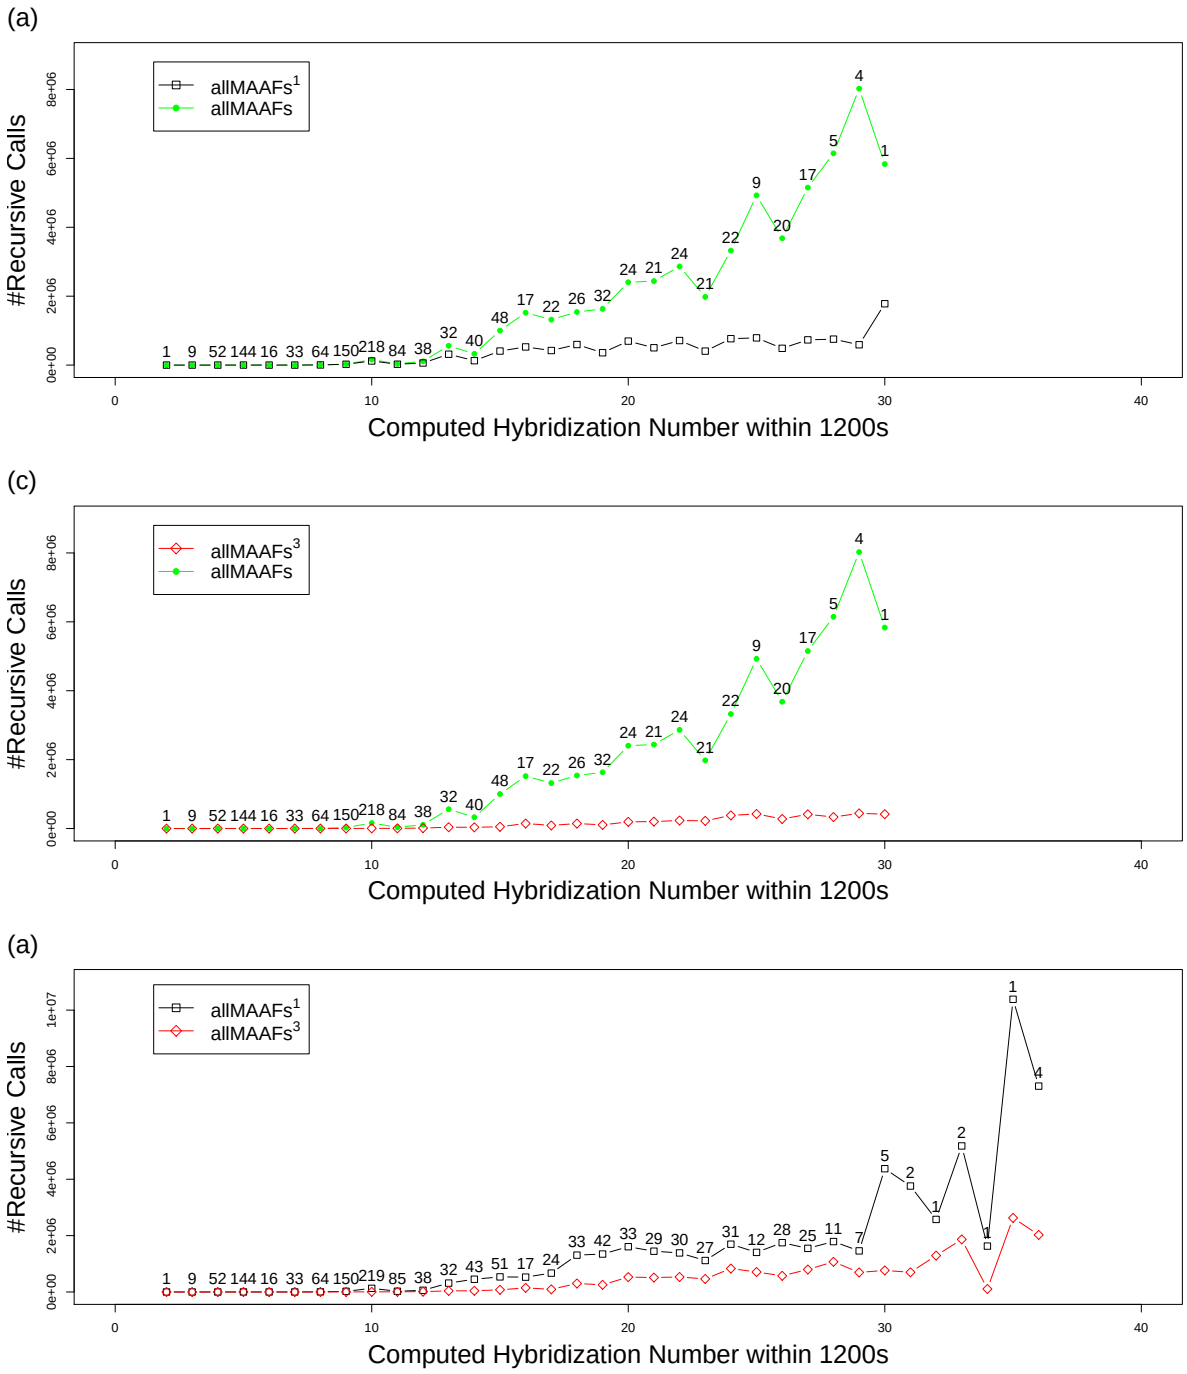

Fig. 9. Comparisons between the mean average number of recursive calls in terms of the hybridization number of (a) allMAAFs and allMAAFs^1 as well as (c) allMAAFs and allMAAFs^2. For each hybridization number the corresponding number of tree pairs is given that could be computed within the time limit of 20 minutes by the two respective algorithms and, thus, contributed to the mean average number of recursive calls denoted by the y-axis.

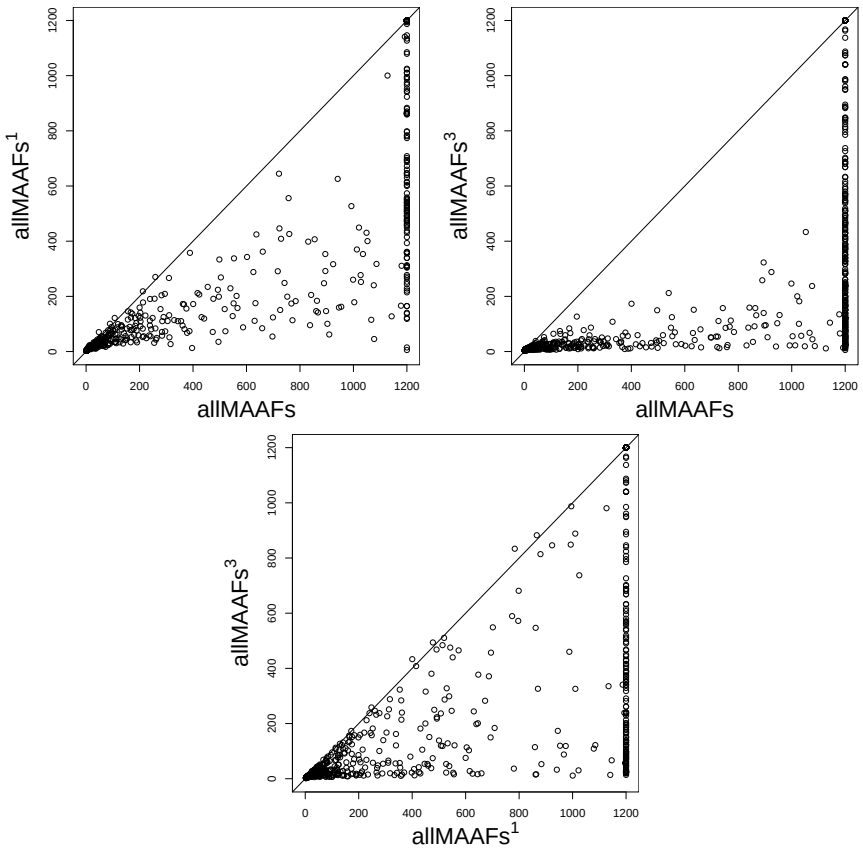

Fig. 10. Scatter-plots comparing the runtimes produced by the three different algorithms when applying our synthetic dataset. Each dot refers to a tree set of this data set whose corresponding x- and y-value indicates the runtime attained by the respective two algorithms.

20 minutes. The plots clearly show that our modified algorithms outperform the original algorithm whereat the modified algorithm allMAAFs^2 is more efficient than the modified algorithm allMAAFs^1.

7 Running time of allMAAFs^1 and allMAAFs^2

The theoretical worst-case runtime of allMAAFs^1 is still the same as the one of the original algorithm allMAAFs, which is $O(3^{|\mathcal{X}|}|\mathcal{X}|)$ as stated in the work of Scornavacca *et al.* [12, Theorem 3].

Theorem 4. *Let T_1 and T_2 be two rooted phylogenetic $\mathcal{X}$-trees and let k be an integer. The algorithm* allMAAFs$^1(T_1, T_2, T_1, \{T_2\}, k)$ *has a worst-case running time of $O(3^{|\mathcal{X}|}|\mathcal{X}|)$.*

Proof. Let $\mathcal{F} = \{F_\rho, F_1, F_2, \ldots, F_{k-1}\}$ be an agreement forest for T_1 and T_2 of size k. To obtain $\mathcal{F}$ from T_2, one obviously has to cut $k-1$ edges. Moreover, in order to reduce the size of leaf set $\mathcal{X}$ of R to 1, to each component F_i in $\mathcal{F}$ we have to contract exactly $|\mathcal{L}(F_i)| - 1$ cherries. Consequently, at most $|\mathcal{X}|$ cherry contractions have to be performed in total. Thus, our algorithm has to perform at most $k + |\mathcal{X}| = O(|\mathcal{X}|)$ recursive calls for the computation of $\mathcal{F}$. Now, as one of these recursive calls can at least branch into three further recursive calls, $O(3^{|\mathcal{X}|})$ is an upper bound for the total number of recursive calls that are performed throughout the whole algorithm. Moreover, as each operation that is performed during a recursive call can be performed in $O(|\mathcal{X}|)$ time, the worse-case running time of both algorithms is $O(3^{|\mathcal{X}|}|\mathcal{X}|)$.

When considering the theoretical worst-case runtime of allMAAFs^2, we have to take the running time of a refinement step into account.

Theorem 5. *Let T_1 and T_2 be two rooted phylogenetic $\mathcal{X}$-trees and let k be an integer. The algorithm* allMAAFs$^2(T_1, T_2, T_1, \{T_2\}, k)$ *has a worst-case running time of $O(3^{|\mathcal{X}|}4^k|\mathcal{X}|)$.*

Proof. As stated in Theorem 4, the algorithm has to conduct $O(3^{|\mathcal{X}|})$ recursive calls. Potentially, for each of these recursive calls we have to apply a refinement step whose theoretical running time is stated with $O(4^k|\mathcal{X}|)$ [14]. Moreover, as all other operations that are performed during a recursive call can be performed in constant time, the worse-case running time of the algorithm is $O(3^{|\mathcal{X}|}4^k|\mathcal{X}|)$.

Even though both presented modifications do not improve the theoretical worst-case running time, our simulation study indicates that in practice these modifications are significantly faster than the original algorithm allMAAFs. Regarding our first modification allMAAFs^1, this is simply due to the fact that the number of computational paths arising from processing a contradicting cherry is reduced. More precisely, if there is a computational path calculating a maximum acyclic agreement forest in which the set of pendant edges E_B for a particular cherry is cut, the original algorithms produces $|E_B| - 1$ redundant recursive calls.

Due to the refinement steps that are performed at the end of each recursive call, our third modification allMAAFs^2 always has to start only one recursive call for processing a common cherry. This refinement step, however, has to be conducted for each maximum agreement forest $\mathcal{F}$ in order to transform $\mathcal{F}$ into all maximum acyclic agreement forests. Nevertheless, as indicated by our simulation study, this post processing step is still more efficient than always running two additional recursive calls when processing a common cherry (as it is the case for the original algorithm allMAAFs).

In conclusion, due to the significant speedup that can be obtained from applying both presented modifications, we feel certain that this work describes a noticeable improvement to the original algorithm allMAAFs, which will make this algorithm accessible for a wider range of biological real-world applications.

8 Conclusion

An important approach to study reticulate evolution is the reconciliation of incongruent gene trees into certain kinds of phylogenetic networks, so-called hybridization networks. From a biological point of view, it is important to calculate not only one but *all* of those networks as, once calculated all of these networks, one can then apply certain filtering techniques testing certain hypothesis. As already discussed in Scornavacca *et. al.* [12], the recently published algorithm allMAAFs calculating all maximum acyclic agreement forests for two rooted binary phylogenetic $\mathcal{X}$-trees, is considered to be an important step to achieve this goal.

In this paper, we have presented the two modifications allMAAFs^1 and allMAAFs^2 of the algorithm allMAAFs still calculating all maximum acyclic agreement forests for both input trees as shown by formal proofs. Moreover, as we have additionally indicated by a simulation study that both modifications significantly improve the practical running time of the original algorithm, we feel certain that this work will facilitate the computation of hybridization networks for larger input trees and, thus, makes this approach accessible to a wider range of biological problems.

Acknowledgements

We gratefully acknowledge Simone Linz for useful discussions.

Bibliography

[1] Albrecht, B.; Scornavacca, C.; Cenci, A.; Huson, D. H.: Fast computation of minimum hybridization networks. *Bioinformatics*, 28(2): 191–197, 2011.

[2] Albrecht, B.: Computing Hybridization Networks for Multiple Rooted Binary Phylogenetic Trees by Maximum Acyclic Agreement Forests. arXiv:1408.3044, 2014.

[3] Baroni, M.; Gruenewald, S.; Moulton, V.; Semple, C.: Bounding the number of hybridisation events for a consisten evolutionary history. *Mathematical Biology* 51: 171–182, 2005.

[4] Baroni, M.; Semple, C.; Steel, M.: Hybrids in real time. *Systematic Biology*, 55: 46–56, 2006.

[5] Bordewich, M.; Semple, C.: Computing the minimum number of hybridization events for a consistent evolutionary history. *Discrete Applied Mathematics*, 155: 914–928, 2007.

[6] Bordewich, M.; Semple, C.: Computing the Hybridization Number of Two Phylogenetic Trees Is Fixed-Parameter Tractable. *IEEE/ACM Trans. Comput. Biol. Bioinformatics*, 4(3): 458–466, 2007.

[7] Chen, Z. Z.; Wang, L.: HybridNET: a tool for constructing hybridization networks. *Bioinformatics*, 26: 2912–2913, 2010.

[8] Collins, J.; Linz, S.; Semple, C: Quantifying hybridization in realistic time. *Journal of Computational Biology*, 18(10): 1305–1318, 2011.

[9] Huson, D.; Rupp, R.; Scornavacca, C.: Phylogenetic Networks: Concepts, Algorithms and Applications. *Cambridge University Press*, 2011.

[10] Huson, D.; Richter, D.; Rausch C.; Dezulian T.; Franz M.; Rupp R.: Dendroscope: An interactive viewer for large phylogenetic trees. *BMC Bioinformatics*, 22;8(1):460, 2007.

[11] Mallet,J.: Hybrid speciation. *Nature*, **446**, 279-283, 2007.

[12] Scornavacca, C.; Simone, L.; Albrecht, B.: A first step towards computing all hybridization networks for two rooted binary phylogenetic trees. *Journal of Computational Biology*, 2012.

[13] Whidden, C; Zeh, N.: A unifying view on approximation and fpt of agreement forests. *Algorithms in Bioinformatics*, LNCS Vol. 5724, 390–402, 2009.

[14] Whidden, C.; Beiko, R.; Zeh, N.: Fixed-Parameter and Approximation Algorithms for Maximum Agreement Forests. arXiv:1108.2664, 2011.